Model risk of contingent claims

Nils Detering and Natalie Packham

November 30, 2014

Abstract

Paralleling regulatory developments, we devise value-at-risk and expected shortfall type risk measures for the potential losses arising from using misspecified models when pricing and hedging contingent claims. Essentially, losses from model risk correspond to losses realized on a perfectly hedged position. Model uncertainty is expressed by a set of pricing models, relative to which potential losses are determined. Using market data, a unified loss distribution is attained by weighing models according to a relative likelihood criterion. Examples demonstrate the magnitude of model risk and corresponding capital buffers necessary to sufficiently protect trading book positions against unexpected losses from model risk.

JEL Classification: G32, G13

Keywords: Model risk, parameter uncertainty, hedge error, value-at-risk, expected shortfall

1 Introduction

Banks and financial institutions closely monitor the market risk associated with their trading activities both for internal risk management and regulatory purposes. Recent updates in banking regulatory frameworks (see BIS, 2011; Federal Reserve, 2011; EBA, 2012) additionally require financial institutions to assess the *model risk* associated with their trading activities, that is, the risk of losses due to using a misspecified model for pricing and hedging securities.[1][2]

Corresponding author: Nils Detering, Department of Mathematics, University of Munich, Theresienstr. 39, 80333 Munich, Germany. Email: nils.detering@math.lmu.de, Phone: +49 89 2180-4579, Fax: +49 89 2180-4452. Natalie Packham, Department of Finance, Frankfurt School of Finance & Management, Sonnemannstr. 9-11, 60314 Frankfurt am Main, Germany. Email: n.packham@fs.de, Phone: +49 69 154008-723, Fax: +49 69 154008-4723.

This work was supported by the Frankfurt Institute for Risk Management and Regulation and by the Europlace Institute of Finance.

We would like to thank Nicole Branger, Pierre Contencin, Jean-Marc Eber, Patrick Hénaff, Jean-Paul Laurent, Pierre-Emmanuel Lévy dit Véhel, Claude Martini, Sébastien Ray, Wolfgang M. Schmidt, Radu Tunaru, participants at the Annual Meeting of the German Finance Association 2013, Quant Congress USA 2013, the 2013 Asian Meeting of the Econometric Society, the 3rd International Conference of the Financial Engineering and Banking Society, the stochastic analysis seminar at the University of Oslo, the model validation group at DekaBank, the Finance research colloquium at Manchester Business School, the Seminar on Model Validation at Institut Louis Bachelier and the Research Seminar of Kent Centre for Finance at the University of Kent for helpful comments and discussions.

[1]For example, BIS (2011) states: "For complex products including, but not limited to, securitisation exposures and n-th-to-default credit derivatives, banks must explicitly assess the need for valuation adjustments to reflect two forms of model risk: the model risk associated with using a possibly incorrect valuation methodology; and the risk associated with using unobservable (and possibly incorrect) calibration parameters in the valuation model." Federal Reserve (2011) states that "model risk should be managed like other types of risk" and that "banks should identify the sources of [model] risk and assess the magnitude".
EBA (2012) states: "Institutions should include the impact of valuation model risk when assessing the prudent value of its balance sheet. [..] Where possible an institution should quantify model risk by comparing the valuations produced from the full spectrum of modelling and calibration approaches."

[2]Throughout, we use the terms "model uncertainty" and "model risk" interchangeably. Uncertainty in the sense of Knight (1921) expresses that beyond the uncertainty associated with the outcome of an event, there

This is a consequence of both increased use of and exposure to models over the last decades as well as the recent experience during the subprime crisis of severe losses from supposedly hedged positions.

The purpose of this paper is to devise risk measures for quantifying potential losses from model risk. To this end, we link model risk to the way a contingent claim can be hedged: In a complete and frictionless market, market risk on a position can be eliminated by hedging, and consequently any observed profits and losses (P&L) on a perfectly hedged position are due to hedging in a misspecifed model. Model risk therefore exists if a position can be hedged only with a model-dependent hedging strategy. Accordingly, risk measures on the "residual" loss of a perfectly hedged position serve as measures of model risk or model uncertainty. In an incomplete market, a clear distinction into market risk and model risk may not be possible, but measuring the "residual" risk, which embeds the model risk, is still possible.

An accurate assessment of model risk when trading contingent claims is important for several reasons: First, assessing the potential losses associated with a claim from using a model for pricing adds to the proper understanding of risks in the trading book beyond market risk. Second, revealing potentially high losses from model uncertainty inherent in a position can prevent both unintentional risk-taking and risk-related incentive conflicts. Third, an adequate assessment of model risk is suitable for deriving capital requirements against unexpected losses from model risk. Section 2 contains an empirical example to demonstrate the scale of model risk involved even in plain vanilla derivatives.

We develop value-at-risk and expected shortfall type measures for model risk, allowing for a direct comparison of model risk with other risk types such as market risk, credit risk and operational risk. This approach coincides with current discussions on regulatory requirements for unexpected losses from model risk (e.g., EBA, 2012).[3] Setting risk limits or implementing capital requirements for model risk expressed in the same units as market risk reduces potential incentives for entering overly model risky positions that appear risk-free when neglecting the model dependence inherent in the hedging strategy. From a regulatory point of view, such risk measures can even prevent systematic mispricing and risk misconceptions of product innovations, and as such reduce systemic risk in the financial system.

Deriving distribution-based measures of model uncertainty, such as value-at-risk and expected shortfall, requires estimating the distribution of losses from hedging in a model-dependent way. Model uncertainty is expressed via a set consisting of alternative models for the asset price dynamics. This gives rise to a set of equivalent martingale measures that are suitable for pricing and hedging. We shall require models to calibrate sufficiently well to liquidly traded options by penalizing models with high calibration error. Since liquidly tradeable options are available only for selected maturities and strikes, both the model and its parameters are not uniquely specified, and it is essentially this uncertainty that should be captured by the set of models.

In a first step, we consider the loss over a pre-specified horizon that arises when hedging in one model – the model used for pricing and hedging – relative to one other model. In a second step, the losses relative to each of the models from the model set are probability-weighted, yielding a unified loss distribution. We demonstrate how probability weights can be derived via techniques from model selection using the Akaike Information Criterion (AIC) (see e.g., Akaike, 1973; Burnham and Anderson, 2002, 2004), which in our case trades off calibration error and model complexity. This gives rise to a market information based estimate of the loss distribution, which in turn is the basis for defining model risk measures associated with a particular hedging strategy. To determine probability weights, one could further include historical information,

is uncertainty associated with the probabilistic behaviour of the event, and it is this latter uncertainty that "model uncertainty" refers to. "Model risk" is concerned with quantifying and measuring the degree of model uncertainty. It is the term prevalent in the finance industry.

[3]EBA (2012) proposes the calculation of a so-called *Additional valuation adjustment (AVA)*, which is the difference in the prudent value and the fair value of a financial product, with the prudent value accounting for unexpected losses at an e.g. 95% confidence level due to model risk amongst other things.

for example by backtesting the hedge quality of each model. This is beyond the scope of this paper.

The risk measures defined entail that static hedging decreases model uncertainty when compared to dynamic hedging. The possibility of (partial) static hedging of a claim is intrinsically connected to model uncertainty (Jarrow, 2012; Ahn and Wilmott, 2008). This is most notably demonstrated by the so-called Breeden-Litzenberger formula, cf. Breeden and Litzenberger (1978); Dupire (1994); Carr and Madan (1998).

We restrict the pricing model to represent a complete market, which allows for a clear differentiation of P&L into market risk and model risk. This distinction becomes blurred in incomplete markets in which perfect hedging strategies neutralising the market risk fail to exist. The main ideas of the approach can still be applied, as one can estimate the loss distribution of the "unhedged" part of the overall P&L, but this will result in some overlap of market risk and model risk. An extension to incomplete markets is treated in Detering and Packham (2014).

From a practical perspective, requiring completeness is not an overly strong restriction: provided there are (few) liquidly traded options that can be used for hedging, completeness is achieved for many diffusion-type stochastic volatility models; for example, in the Heston stochastic volatility model, one liquidly traded option is sufficient to achieve market completeness.

Parallel to the development of models for valuing and hedging contingent claims, there has always been a natural interest in understanding the risks associated with employing models, (e.g. Galai, 1977; Merton, Scholes, and Gladstein, 1978, 1982; Figlewski, 1998); overviews are given in e.g. Derman (1996); Crouhy, Galai, and Mark (1998); Hénaff and Martini (2011); Morini (2011). A large number of papers analyse the variation in prices and hedging strategies across different models, typically for certain classes of models or payoffs (e.g. Carr and Madan, 1998; Hull and Suo, 2002; Nalholm and Poulsen, 2006; Branger, Krautheim, Schlag, and Seeger, 2012).

Only few papers, however, are concerned with determining capital requirements for model risk. Kerkhof, Melenberg, and Schumacher (2010) directly integrate uncertainty about the model into value-at-risk for market risk. Elices and Giménez (2013) determine and compare losses from hedging in a Black-Scholes and volga-vanna model assuming that the market follows a Heston model. (Bannör and Scherer, 2013) derive an expected shortfall-type risk measure from the uncertainty associated with parameter estimation. (Glasserman and Xu, 2014) determine model risk in a worst-case scenario across portfolio strategies. A number of recent papers – Boucher, Danielsson, Kouontchou, and Maillet (2014); Breuer and Csiszar (2014); Boyarchenko, Cerrato, Crosby, and Hodges (2014); Hénaff and Martini (2011); Cohort, Vehel, and Patras (2013) – treat the issues of model uncertainty and model risk, but do not calculate a capital risk charge as a result of measuring unexpected losses.

In the light of the recent credit subprime crisis, losses associated with credit securisations (e.g., van Deventer, 2008; Heitfield, 2008; Jorion, 2009; Ascheberg, Bick, and Kraft, 2013) demonstrate the importance of accounting for model risk, for example in terms of adjusting profitability and building reserves for model risk. Particularly striking examples were so-called *Constant Proportion Debt Obligations (CPDOs)*, Gordy and Willemann (2012); Cont and Jessen (2012), and leveraged credit securities, Morini (2011); Packham, Schloegl, and Schmidt (2013), whose valuation and dynamics are extremely sensitive to model assumptions.

The paper is structured as follows: To motivate the analysis further, we introduce an empirical example of P&L generated from misspecified hedging in Section 2. Section 3 contains the market setup; further, the loss process from hedging is introduced. The distribution of losses from model risk relative to a set expressing the model uncertainty is derived in Section 4. One method of defining probability weights on the models is via the Akaike Information Criterion. Model risk measures suitable for defining capital requirements are defined in Section 5. Furthermore, we show that these measures fulfill the axioms for measures of model uncertainty devised by Cont (2006). Section 6 contains several examples and Section 7 concludes.

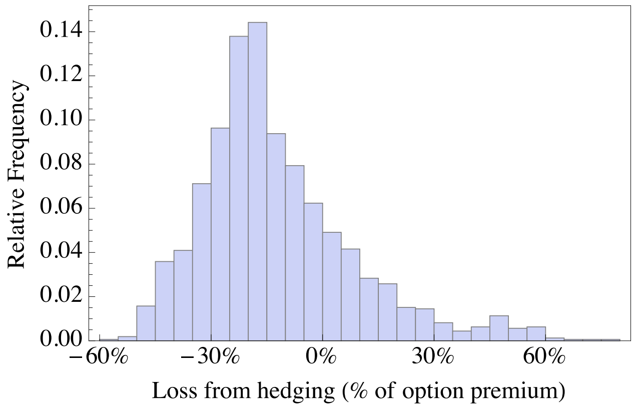

Figure 1: Empirical P&L distribution from hedging 3-month at-the-money call options on the DAX index in a Black Scholes model with rebalancing every five minutes. P&L shown as a percentage of the option premium at inception (Data source: Bloomberg, daily data from October 2006 until April 2013).

2 Empirical example

To motivate our analysis further, we study the empirical loss distribution from delta-hedging a call option on the DAX index. This is a sufficiently liquid and mature market to warrant that the P&L observed from a hedged position is indeed due to model misspecification. Implied volatilities are given by the DAX volatility index VDAX-NEW (Bloomberg ticker VDAX3M), which is a measure of implied volatility derived from traded 3-month DAX options on the Eurex derivatives exchange. As a risk-free interest rate we take the 3-month LIBOR (Bloomberg ticker EE0003M). The time horizon is 24 October 2006 until 22 April 2013.

In our example, on each trading day a 3-month at-the-money call option is written. Each option is delta-hedged in the Black-Scholes model using a self-financing replicating strategy with the DAX future. The implied volatility entering the formula for the Black-Scholes Delta is the VDAX-NEW volatility from the day of inception. The re-balancing of the hedge portfolio is done every five minutes to ensure that the discretization error is negligible. The P&L at option expiry is then taken as the error from hedging in a misspecified model (Black-Scholes in this case). The empirical P&L distribution of the 1,588 realizations obtained is shown in Figure 1. The expected profit is consistent with the well-known risk premium for volatility (e.g. Carr and Wu, 2009). Overall, however, the realised variation is large and we observe a 28% loss of the option premium at a 95% confidence level and a 54% loss of the option premium at the 99% confidence level.

Of course, in practice, traders are unlikely to use the Black-Scholes model with a constant volatility over time for hedging and, in addition, traders do not hedge individual options but manage trading books. Nonetheless, the example is instructive in highlighting that losses from model misspecification can be substantial and as such cannot be ignored. Also, updating an option's implied volatility, which corresponds to re-calibrating the model over time, did not yield a significant improvement.

In the light of this example, consider briefly the price range measure, which is a popular and simple measure of model risk used in practice (e.g., Schoutens, Simons, and Tistaert, 2004; Cont, 2006). Using several pricing models, the price range of a claim is just the difference between the greatest and smallest prices. The example demonstrates that unexpected losses from model misspecification must take into account misspecified dynamics. Such losses are not captured by a measure based on the price range, and, as such, a capital requirement based on the price range would fail to act as a capital buffer against unexpected losses from model risk. In the extreme case, the price of a payoff could be equal across all models, with the hedging strategies differing across models. In this case the price range measure is zero, but the losses from model risk can be substantial. Such an example is given in Appendix A.

3 Market setup

3.1 Market and model setup

We begin with a standard market setup under model certainty. On a a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ endowed with a filtration $(\mathcal{F}_t)_{t\geq 0}$ satisfying the "usual hypotheses" are defined adapted asset price processes $(S_t^j)_{t\geq 0}$, $j = 0, \ldots, d$. The asset with price process S^0 represents the money market account, whereas $S^1, \ldots, S^d$ are risky assets. All prices are discounted, that is, expressed in units of the money market account. Further, there exists an equivalent martingale measure $\mathbb{Q}$, under which all asset price processes are $\mathbb{Q}$-martingales, making the market arbitrage-free. We fix a time horizon T and we consider claims with $\mathcal{F}_T$-measurable integrable payoff.

In addition to the risky assets $S = (S^1, \ldots, S^d)$, there may be tradeable options written on S, with $\mathcal{F}_T$-measurable payoff and with observable market prices at time 0, so-called *benchmark instruments*. Their $\mathcal{F}_T$-measurable payoffs are denoted by $(H_i)_{i\in I}$, and their observed market prices by C_i^*, $i \in I$, or by $[C_i^{\text{bid}}, C_i^{\text{ask}}]$, $i \in I$, if no unique price is available. These benchmark instruments can be used for static hedging, potentially reducing a claim's model risk.[4] Semi-static hedging in benchmark instruments, which is important in practice, will be discussed in Section 6.3.

A trading strategy (or portfolio) is a predictable process $\Phi = (\phi^0, \ldots, \phi^d, u_1, \ldots, u_I)$, where $\phi^j = (\phi_t^j)_{t\geq 0}$ denotes the holdings in asset j and $u_i \in \mathbb{R}$ denotes the static holding of benchmark instrument i. The time-t value of the portfolio is

$$V_t(\Phi) = \sum_{j=0}^{d} \phi_t^j S_t^j + \sum_{i=1}^{I} u_i H_t^i, \tag{1}$$

with H_t^i, $i = 1, \ldots, I$, the time-t prices of the benchmark instruments. To rule out arbitrage opportunities we require that Φ is admissible. Further, Φ is assumed to be self-financing, that is, $\mathrm{d}V_t(\Phi) = \sum_{j=1}^{d} \phi_t^j \,\mathrm{d}S_t^j + \sum_{i=1}^{I} u_i \,\mathrm{d}H_t^i$, $t \geq 0$.

A contingent claim with $\mathcal{F}_T$-measurable payoff X is hedgeable if there exists a replicating strategy, i.e., a self-financing trading strategy Φ such that $V_T(\Phi) = X$. Hedging eliminates any P&L arising from market risk, and, because of the absence of arbitrage opportunities, the claim's price process and the price of the hedging strategy agree for all $0 \leq t \leq T$. In order to clearly distinguish market risk and model risk, we restrict attention to complete markets.

Assumption 1. The market is complete under $\mathbb{Q}$.

Aside from market risk, a stakeholder (trader, hedger, shareholder, regulator, ...) may be concerned about model risk when pricing and hedging a contingent claim. Model risk refers to potential losses from mispricing and mishedging, because model $\mathbb{P}$, resp. $\mathbb{Q}$, may be misspecified. This uncertainty regarding the model $\mathbb{P}$ (resp. the pricing model $\mathbb{Q}$) is captured by a set of

[4]The role of static hedging with benchmark instruments is to reduce model uncertainty (as postulated in the axiomatic setup of Cont (2006), which will be discussed in Section 5.3). The principal idea is that a full static replication of a claim is indeed model-independent. A simple example is put-call parity, where a put option can be statically replicated by a position in the forward and call option. More generally, Carr and Madan (1998) show that any twice differentiable European payoff can be expressed as a static position in bonds, forward contracts and call and put option of arbitrary strikes, which in turn is model-independent.

On the other hand, due to the assumption of market completeness in the underlying assets, dynamic hedging in the benchmark instruments is in principle redundant and therefore does not contribute to reducing model risk, which is why the possibility of dynamic hedging in the benchmark instruments is excluded from the trading strategies considered.

measures, $\mathcal{P}$, for the asset price processes. This incorporates uncertainty about both model type and model parameters. Often, it is sufficient to capture model uncertainty by a set $\mathcal{Q}$ of martingale measures (e.g. Cont, 2006; Denis and Martini, 2006), for example when model risk is measured by price differences, where only martingale measures matter. In practice, however, since risk is measured under the objective measure, it may be necessary to capture the uncertainty by the set $\mathcal{P}$ rather than $\mathcal{Q}$. For every measure in $\mathcal{P}$ we assume existence of an equivalent martingale measure, so that $\mathcal{Q}$ consists of equivalent martingale measures relative to $\mathcal{P}$.[5]

Working on a set of measures requires further conditions, in particular, as the measures in $\mathcal{P}$ need not be absolutely continuous with respect to $\mathbb{P}$. More specifically, the asset price processes must be consistent under all measures in $\mathcal{P}$, and specifying trading strategies requires the notion of a stochastic integral with respect to $\mathcal{P}$.

Assumption 2. For any admissible and predictable trading strategy ϕ, there exists a version of the stochastic integral $\int_0^t \phi \, \mathrm{d}S$, such that, for all $\hat{\mathbb{P}} \in \mathcal{P}$, the integral coincides $\hat{\mathbb{P}}$–a.s. with the usual probabilistic construction and $\int_0^t \phi \, \mathrm{d}S$ is $\mathcal{F}_t$-measurable.

In case the models in $\mathcal{P}$ are diffusion processes, Soner, Touzi, and Zhang (2011) develop the required tools from stochastic analysis, such as existence of a stochastic integral, martingale representation, etc. Although this restricts the joint occurrence of certain probability measures, it does not exclude any particular measure. For our purposes, this limitation does not play a role, as the primary interest lies in choosing a rich set of possible models to cover the model uncertainty.

One way to make the conditions precise is to adopt the setting of Soner et al. (2011), but it should be noted that they can be accomplished in other settings as well. In this special setup, the conditions are as follows:

(i) The filtration is completed in the sense of Definition 2.2 of Soner et al. (2011), implying that it is right-continuous, but not necessarily complete under each $\hat{\mathbb{P}} \in \mathcal{P}$.
(ii) The set $\mathcal{P}$ fulfills the following conditions:
  - $S = (S^0, \ldots, S^d)$ is a diffusion for every $\hat{\mathbb{P}} \in \mathcal{P}$ and an *aggregator*, that is, $S = S^{\hat{\mathbb{P}}}$ $\hat{\mathbb{P}}$–a.s. for every $\hat{\mathbb{P}} \in \mathcal{P}$, with $S^{\hat{\mathbb{P}}}$ the discounted price process under $\hat{\mathbb{P}}$ (and a square-integrable $\hat{\mathbb{Q}}$-martingale under the corresponding equivalent martingale measure(s) $\hat{\mathbb{Q}} \in \mathcal{Q}$).
  - $\mathcal{P}$ fulfills the separability and consistency conditions of Definition 4.8 and Theorem 5.1 of Soner et al. (2011).

The set of contingent claims under consideration is given by

$$\mathcal{C} = \left\{ X \in \mathcal{F}_T \,\middle|\, \mathbb{E}^{\mathbb{Q}}[|X|] < \infty \right\},$$

and the set of trading strategies considered is

$$\mathcal{S} = \Bigg\{ \phi \,\Bigg|\, \phi \text{ admissible, self-financing, } (\mathcal{F}_t)_{t \leq T}\text{-predictable} \\ \text{and } \mathbb{E}^{\mathbb{Q}}\left[\int_0^T (\phi^j)^2 \, \mathrm{d}[S^j, S^j] \right] < \infty, j = 0, \ldots, d \Bigg\}.$$

[5] If $\mathcal{P}$ contains incomplete market models, then there is some flexibility in the specification of $\mathcal{Q}$. For example, it may be composed of *all* equivalent martingale measures relative to $\mathcal{P}$, or it may contain one "representative" martingale measure for each measure in $\mathcal{P}$.

3.2 Loss process

Recall that a trading strategy (or portfolio) is a predictable self-financing process $\Phi = (\phi^0, \dots, \phi^d, u_1, \dots, u_I)$, where $\phi^j = (\phi^j_t)_{t\geq 0} \in \mathcal{S}$ denotes the holdings in asset j and $u_i \in \mathbb{R}$ denotes the static holding of benchmark instrument i. Consider a short position in a claim $X \in \mathcal{C}$ and a hedging strategy $\Phi = (\phi, u_1, \dots, u_I)$. We assume for the moment that the hedging strategy is uniquely given for all $\omega \in \Omega$ (as is the typical case in practice). The time-T loss associated with X is given by

$$L_T(X, \Phi) := -(V_T(\phi) - Y),$$

where $V_T(\phi) = V_T((\phi, 0, \dots, 0))$ is given by Equation (1) and $Y = X - \sum_{i=1}^I u_i H_i$. In other words, $L_T(X, \Phi)$ measures loss associated with the dynamically hedged part of the claim. If $\mathbb{Q}$ calibrates to the market prices of the benchmark instruments, i.e., $\mathbb{E}[H_i] = C_i^\star$, $i = 1, \dots, I$, then $L_T(X, \Phi) = -(V_T(\Phi) - X)$ (for notational convenience, we associate with $\mathbb{E}$ the expectation under the pricing measure $\mathbb{Q}$). Otherwise, if $\mathbb{Q}$ fails to calibrate perfectly to the benchmark instruments, then entering into the static positions produces an initial P&L, which, despite being generated by (obvious) model misspecification, is considered a trading cost and excluded from L_T: First, it is a sunk cost and there is no uncertainty associated with this P&L, so one does not need to provision for it. Second, there is no further P&L associated with the statically hedged part of the position, provided the claim and the statically hedged part are held until maturity.

In general, ϕ will be defined only $\mathbb{Q}$–a.s., and one must be explicit in specifying the version to be used when dealing with models that are not absolutely continuous with respect to $\mathbb{Q}$. Likewise, extending the loss variable L_T to a loss process

$$L_t := L_t(X, \Phi) = -(V_t(\phi) - \mathbb{E}[Y|\mathcal{F}_t]), \quad 0 \leq t \leq T,$$

with Φ the replicating strategy under $\mathbb{Q}$, requires that the version of the time-t price $\mathbb{E}[Y|\mathcal{F}_t]$ be explicitly specified. As a minimal requirement, since $\mathcal{P}$ (resp. $\mathcal{Q}$) expresses the model uncertainty when employing $\mathbb{Q}$ for pricing and hedging, it must not be involved in the choice of the respective version representing the pricing and hedging strategies.[6] For example, from a risk-control perspective, this would reflect that model risk is not calculated at the trading desk, but at an independent risk management unit.

Furthermore, a minimal requirement that will be important for defining meaningful risk measures is that linearity on the versions chosen is preserved for all $\omega \in \Omega$ and not only $\mathbb{Q}$–a.s.. Recall that $\mathbb{E} \equiv \mathbb{E}^{\mathbb{Q}}$. To simplify notation we shall assume that $\mathcal{F}_0$ is trivial and often simply write $\mathbb{E}[Y]$ instead of $\mathbb{E}[Y|\mathcal{F}_0]$.

Assumption 3.

(i) For any claim $Y \in \mathcal{C}$, a unique ($\forall \omega \in \Omega$) price process $(\mathbb{E}[Y|\mathcal{F}_t])_{t\geq T}$ with $\mathbb{E}[Y|\mathcal{F}_T] = Y$ and a unique replicating strategy $(\phi_t(Y))_{t\geq T}$ are chosen, irrespective of the measures contained in $\mathcal{P}$.

(ii) If the trading strategy $\phi(Y)$ is a deterministic function of time $\mathbb{Q}$–a.s., then the deterministic version is chosen for all $\omega \in \Omega$, and $\mathbb{E}[Y|\mathcal{F}_t] = \sum_i \int_0^t \phi^i(Y)\, \mathrm{d}S^i$, $t \leq T$, provided the right-hand side exists.

(iii) For any two claims $Y_1, Y_2 \in \mathcal{C}$

$$\mathbb{E}[aY_1 + bY_2|\mathcal{F}_t] = a\mathbb{E}[Y_1|\mathcal{F}_t] + b\mathbb{E}[Y_2|\mathcal{F}_t], \quad a, b \in \mathbb{R}$$

[6] Since, suppose for example, that $\mathcal{Q} = \{\mathbb{Q}, \overline{\mathbb{Q}}\}$, and $\mathbb{Q}$ and $\overline{\mathbb{Q}}$ are singular measures. Then, knowledge of $\mathcal{Q}$ could be used to choose a trading strategy replicating the claim under both $\mathbb{Q}$ and $\overline{\mathbb{Q}}$, eliminating any model risk and thus rendering $\mathcal{Q}$ unsuitable for expressing model uncertainty. But this is impracticable and therefore needs to be ruled out.

and

$$a\phi(Y_1) + b\phi(Y_2)) = \phi(aY_1 + bY_2), \quad a, b \in \mathbb{R}.$$

In practice, a claim's price and hedging strategy are typically determined from the current asset prices and the pricing model $\mathbb{Q}$. For example, if the claim's price process is Markov with respect to S and can be priced and hedged via a PDE, one would choose the version that solves the related PDE, so there is no ambiguity about the version chosen. Furthermore, Assumption 3 will be automatically fulfilled.

In the following we shall often suppress the dependency of $\phi(Y)$ on Y where it is clear from the context. For notational convenience, we shall stick to the notation of the conditional expectation $\mathbb{E}[Y|\mathcal{F}_t]$ (which is defined only $\mathbb{Q}$–a.s.), but we shall always assume that $\mathbb{E}[Y|\mathcal{F}_t]$ corresponds to a version fulfilling Assumption 3.

Definition 4. Let $X \in \mathcal{C}$ and $\Phi = (\phi, u_1, \ldots, u_I)$ with $Y = X - \sum_{i=1}^I u_i H_i$ and $\phi = \phi(Y)$. The *loss process* associated with a short position in X and the trading strategy Φ is given by

$$\begin{aligned} L_t = L_t(X, \Phi) &= -(V_t(\phi) - \mathbb{E}[Y|\mathcal{F}_t]) \\ &= -\big(V_0 + \sum_{j=1}^d \textstyle\int_0^t \phi^j \, \mathrm{d}S^j - \mathbb{E}[Y|\mathcal{F}_t]\big), \quad 0 \le t \le T, \end{aligned}$$

with $V_0 = \mathbb{E}[Y]$.

In other words, L_t is just minus the P&L observed until time t.

4 The distribution of losses from model risk

In a frictionless market, any P&L observed on a perfectly hedged position is due to hedging in a misspecified model. Assumption 1, which postulates a complete market, justifies the term "model risk". The price process of Y when pricing according to model $\mathbb{Q}$ is given by $\mathbb{E}[Y|\mathcal{F}_t]$ $\mathbb{Q}$–a.s. and by definition any replicating strategy ϕ is such that $L_t = 0$ $\mathbb{Q}$–a.s., resp. $L_t = 0$ $\mathbb{P}$–a.s., $t \le T$. On the other hand, $\hat{\mathbb{P}}(L_t = 0) < 1$, for some $\hat{\mathbb{P}} \in \mathcal{P}$, expresses that ϕ fails to replicate Y under $\hat{\mathbb{P}}$.

A model-free hedging strategy is defined as follows:

Definition 5. The trading strategy $\Phi = ((\phi_t)_{0 \le s \le T}, u_1, \ldots, u_I)$ is a *model-free* or *model-independent* hedging strategy for claim X with respect to $\mathcal{P}$, if $L_t = 0$, $0 \le t \le T$, $\hat{\mathbb{P}}$–a.s., for all $\hat{\mathbb{P}} \in \mathcal{P}$.

Because the set $\mathcal{Q}$ consists of equivalent martingale measures of the measures contained in $\mathcal{P}$, model-independent strategies may equivalently be defined relative to $\mathcal{Q}$ instead of $\mathcal{P}$.

4.1 Loss from hedging relative to one model

In a first step, we investigate the loss from hedging when the market evolves according to $\hat{\mathbb{Q}} \in \mathcal{Q}$ instead of $\mathbb{Q}$.

Proposition 6. *Let $\hat{\mathbb{Q}} \in \mathcal{Q}$ be a complete market and $Y \in \mathcal{C}$ such that*

$$\mathbb{E}^{\hat{\mathbb{Q}}}\left[\int_0^T (\phi^j)^2 \, \mathrm{d}[S^j, S^j]\right] < \infty, j = 0, \ldots, d \tag{2}$$

and $\mathbb{E}^{\hat{\mathbb{Q}}}[(\mathbb{E}[Y|\mathcal{F}_t])^2] < \infty$, where $\phi = \phi(Y)$ and $\mathbb{E}[Y|\mathcal{F}_t]$ are the particular versions fulfilling Assumption 3 . The loss from hedging under $\mathbb{Q}$ relative to $\hat{\mathbb{Q}}$ can be represented as an integral of the realised path:

$$L_t = -(\mathbb{E}[Y] - \mathbb{E}^{\hat{\mathbb{Q}}}[\mathbb{E}[Y|\mathcal{F}_t]]) - \int_0^t (\phi - \phi^{\hat{\mathbb{Q}},t}) \, \mathrm{d}S, \tag{3}$$

where $\phi^{\hat{\mathbb{Q}},t}$ satisfies $\mathbb{E}[Y|\mathcal{F}_t] = \mathbb{E}^{\hat{\mathbb{Q}}}[\mathbb{E}[Y|\mathcal{F}_t]] + \int_0^t \phi^{\hat{\mathbb{Q}},t}\,\mathrm{d}S$ $\hat{\mathbb{Q}}$–a.s. Furthermore, $\int_0^t(\phi - \phi^{\hat{\mathbb{Q}},t})\,\mathrm{d}S$ is a $\hat{\mathbb{Q}}$-martingale, and

$$\mathbb{E}^{\hat{\mathbb{Q}}}[L_t] = -(\mathbb{E}[Y] - \mathbb{E}^{\hat{\mathbb{Q}}}[\mathbb{E}[Y|\mathcal{F}_t]]). \tag{4}$$

Proof. Since $\mathbb{E}[Y|\mathcal{F}_t]$ is $\mathcal{F}_t$-measurable and square integrable by assumption and since the market under $\hat{\mathbb{Q}}$ is complete, there exists a trading strategy $\phi^{\hat{\mathbb{Q}},t}$ that replicates $\mathbb{E}[Y|\mathcal{F}_t]$, (note that the strategy depends on the time horizon t, since $\mathbb{E}[Y|\mathcal{F}_t]$ is not a martingale under $\hat{\mathbb{Q}}$). By definition of L and $\phi^{\hat{\mathbb{Q}},t}$ we obtain

$$\begin{aligned} L_t &= -(V_t - \mathbb{E}[Y|\mathcal{F}_t]) \\ &= -\left(\mathbb{E}[Y] + \int_0^t \phi\,\mathrm{d}S - \mathbb{E}^{\hat{\mathbb{Q}}}[\mathbb{E}[Y|\mathcal{F}_t]] - \int_0^t \phi^{\hat{\mathbb{Q}},t}\,\mathrm{d}S\right)\ \hat{\mathbb{Q}}\text{–a.s.} \end{aligned} \tag{5}$$

which proves Equation (3). By the assumption on ϕ, Equation (4) is obtained by taking expectation. □

Note that Equation (4) follows directly from the martingale property of the trading gains and especially does not depend on any completeness assumption. Only the representation (3) might be lost if $\hat{\mathbb{Q}}$ is not complete.

Corollary 7. *The total expected loss from hedging claim X, that is $\mathbb{E}^{\hat{\mathbb{Q}}}[L_T]$ plus the initial transaction cost, is just the price difference in the two models, $-(\mathbb{E}[X] - \mathbb{E}^{\hat{\mathbb{Q}}}[X])$.*

Proof. Let $\Phi = ((\phi_t)_{0\le t\le T}, u_1, \dots, u_I)$ be a replicating strategy for X under $\mathbb{Q}$. Then,

$$\begin{aligned} &\mathbb{E}^{\hat{\mathbb{Q}}}[L_T] - \left(\mathbb{E}[\textstyle\sum_i u_iH_i] - \mathbb{E}^{\hat{\mathbb{Q}}}[\sum_i u_iH_i]\right) \\ &\qquad = -(\mathbb{E}[Y] - \mathbb{E}^{\hat{\mathbb{Q}}}[\mathbb{E}[Y|\mathcal{F}_T]]) - \left(\mathbb{E}[\textstyle\sum_i u_iH_i] - \mathbb{E}^{\hat{\mathbb{Q}}}[\sum_i u_iH_i]\right) \\ &\qquad = -(\mathbb{E}[Y] - \mathbb{E}^{\hat{\mathbb{Q}}}[Y]) - \left(\mathbb{E}[\textstyle\sum_i u_iH_i] - \mathbb{E}^{\hat{\mathbb{Q}}}[\sum_i u_iH_i]\right) \\ &\qquad = -(\mathbb{E}[X] - \mathbb{E}^{\hat{\mathbb{Q}}}[X]). \end{aligned}$$

□

With the Corollary one can restate the price range measure from Section 2 (see also Appendix A) as $\mu_{\mathcal{Q}}(X) = \sup_{\hat{\mathbb{Q}},\mathbb{Q}\in\mathcal{Q}} \mathbb{E}^{\hat{\mathbb{Q}}}[L_T^{\mathbb{Q}}]$, where $L_T^{\mathbb{Q}}$ denotes the loss variable from hedging under $\mathbb{Q}$. The capital charge, when pricing under model $\mathbb{Q}$, is given by $\sup_{\hat{\mathbb{Q}}\in\mathcal{Q}} \mathbb{E}^{\hat{\mathbb{Q}}}[L_T^{\mathbb{Q}}]$. In terms of hedging, it measures the worst *expected* loss from hedging (within the model set $\mathcal{Q}$).

4.2 Relation to tracking error

Assuming that Y depends only on the final value of the underlying, that is, $Y = h(S_T)$ for some Borel function h, and under an additional assumption on the volatility process in the misspecified model, the loss from hedging in the misspecified model coincides with the *tracking error* derived in El Karoui, Jeanblanc-Picqué, and Shreve (1998), and L_t admits an explicit expression. In the model used for hedging, $\mathbb{Q}$, the discounted stock price dynamics are given by

$$\mathrm{d}S_t = S_t\sigma(t, S_t)\,\mathrm{d}W_t, \tag{6}$$

where $\sigma : [0,T]\times(0,\infty) \mapsto [0,\infty)$ is continuous and bounded above, and further $(\partial/\partial s)[s\,\sigma(t,s)]$ is Hölder-continuous in (s,t), Lipschitz continuous and bounded in $s \in (0,\infty)$, uniformly in

$t \in [0,T]$, see Hypotheses 5.1 and 6.1 of El Karoui et al. (1998). The stock price dynamics under the risk-neutral market measure $\hat{\mathbb{Q}}$, on the other hand, are given by

$$\mathrm{d}S_t = S_t \hat{\sigma}_t \, \mathrm{d}W_t,$$

with $\hat{\sigma}_t$ $(\mathcal{F}_t)_{t\geq 0}$-adapted and satisfying $\int_0^T \hat{\sigma}^2(t)\,\mathrm{d}t < \infty$ $\hat{\mathbb{Q}}$–a.s.. Since S_t is a Markov process under $\overline{\mathbb{Q}}$ and Y only depends on the final value S_T, the time-t price $\mathbb{E}[Y|\mathcal{F}_t]$ of Y is of the form $v(t, S_t)$ with $v(t,s) \in C([0,T]\times(0,\infty)) \cap C^{1,2}([0,T)\times(0,\infty))$. The time-$t$ loss from hedging is given by, cf. Eq. (6.7) of El Karoui et al. (1998),

$$L_t = -\frac{1}{2}\int_0^t \sigma^2(u, S_u) - \hat{\sigma}^2(u) S_u^2 \frac{\partial^2}{\partial s^2}\Big[v(t, S_t)\Big]\,\mathrm{d}u. \tag{7}$$

If both models have only time-dependent volatility and are calibrated to the same integrated variance at time T, that is, $\int_0^T \hat{\sigma}^2(t)\,\mathrm{d}t = \int_0^T \sigma^2(t, S_t)\,\mathrm{d}t$, then the prices at time 0 of European payoffs depending only on S_T agree in both models. If instead $\hat{\sigma}_t \leq \sigma(t, S_t)$, for Lebesgue-almost all $t \in [0,T]$, and h is a convex function with bounded one-sided derivatives (recall that $Y = h(S_T)$), then $L_t \leq 0$, for $0 \leq t \leq T$, that is, the hedging strategy is a superhedge. Conversely, $\hat{\sigma}_t \geq \sigma(t, S_t)$, $t \in [0,T]$, implies that the hedging strategy is a subhedge, Theorem 6.2 of El Karoui et al. (1998).

Representation (7) allows to characterize claims that can be hedged in a model-free way.

Proposition 8. *Let $\hat{\sigma}_t$ and $\sigma(t, S_t)$ fulfill the properties stated above so that Equation (7) holds and suppose that $\hat{\mathbb{Q}}(\hat{\sigma}_t \neq \sigma(t, S_t),$ for Lebesgue-almost all $t \in [0,T]) = 1$. Then, $\hat{\mathbb{Q}}(L_t = 0) = 1$ for all $t \in [0,T]$ if and only if $Y = h(S_T) = aS_T + b$ with $a, b \in \mathbb{R}$.*

Proof. "$\Rightarrow$" By assumption on the diffusion term $\sigma(t,s)$ it follows that $\mathbb{E}[Y|\mathcal{F}_T] = y(t, S_t)$ for some $y(t,x) \in C([0,T]\times(0,\infty)) \cap C^{1,2}([0,T)\times(0,\infty))$, see page 104 of El Karoui et al. (1998). For $\hat{\mathbb{Q}}(L_t = 0) = 1$ to hold for all $t \in [0,T]$, the integrand in Equation (7) must vanish $\hat{\mathbb{Q}}$-a.s. for Lebesgue-almost all $t \in [0,T]$. Both $\hat{\mathbb{Q}}(\hat{\sigma}_t \neq \sigma(t, S_t),$ for Lebesgue-a.a. $t \in [0,T]) = 1$ and S_t strictly positive imply that $\frac{\partial^2 y(t,S_t)}{\partial x^2} = 0$ $\hat{\mathbb{Q}}$–a.s. for Lebesgue-almost all $t \in [0,T]$. Since $\frac{\partial^2 y(t,S_t)}{\partial x^2}$ is continuous, it follows that $\frac{\partial^2 y(t,S_t)}{\partial x^2} = 0$, for all $t \in [0,T)$. But this implies that $y(t, S_t)$ is linear in S_t for fixed t, more specifically $y(t, S_t) = S_t f(t) + g(t)$, for some continuous functions $f(t)$ and $g(t)$ defined on $[0,T)$. By continuity of $y(t,x)$, $y(T, S_T) = S_T f(T_-) + g(T_-)$, proving the claim with $a = f(T_-)$ and $b = g(T_-)$. To show "$\Leftarrow$" observe that $y(t, S_t) = aS_t + b$ and thus $\frac{\partial^2 y(t,S_t)}{\partial x^2} = 0$. □

4.3 Losses from hedging in a model-dependent way

To extend the loss distribution to include losses relative to all models in $\mathcal{P}$, consider an extended probability space $(\Omega, \mathcal{F}, \overline{\mathbb{P}})$, where $\mathcal{F}$ now incorporates in addition the model uncertainty and $\overline{\mathbb{P}}$ contains information about the degree of uncertainty associated with each model. To make this precise, let $\mathcal{G} \subset \mathcal{F}$ be a σ-algebra such that conditioning on $\mathcal{G}$ eliminates the uncertainty about the measure $\hat{\mathbb{P}} \in \mathcal{P}$. In this setting, the measures that constitute the regular conditional probability with respect to $\mathcal{G}$ are the attainable models and $\mathcal{P}$ corresponds to the measures associated with this regular conditional probability. For the existence and construction of this probability space given the set $\mathcal{P}$, see Appendix B. Without loss of generality we assume existence of a random variable $\theta \in \Theta \subseteq \mathbb{R}$ with $\sigma(\theta) = \mathcal{G}$, so that the elements of $\mathcal{P}$ are indexed by θ, and $\overline{\mathbb{P}}_\theta = \overline{\mathbb{P}}(\,\cdot\,|\sigma(\theta))$, resp. $\overline{\mathbb{P}}_a = \overline{\mathbb{P}}(\cdot\,|\theta = a)$, $a \in \Theta$. Moreover, for $B \in \mathcal{F}$, we have

$$\overline{\mathbb{P}}(B) = \mathbb{E}^{\overline{\mathbb{P}}}[\overline{\mathbb{P}}(B|\sigma(\theta))] = \int_\Omega \overline{\mathbb{P}}(B|\sigma(\theta))\,\mathrm{d}\overline{\mathbb{P}} = \int_\Theta \overline{\mathbb{P}}(B|\theta = a)\,\mu(\mathrm{d}a), \tag{8}$$

where $\mu = \overline{\mathbb{P}} \circ \theta^{-1}$ is the distribution of θ. In this way, model uncertainty is expressed by the unconditional distribution $\overline{\mathbb{P}}$ and model certainty is expressed via the conditional distributions $\overline{\mathbb{P}}(\,\cdot\,|\sigma(\theta))$. This setup endows the set of measures $\mathcal{P}$ with a distribution expressing the degree of uncertainty associated with its elements. One way of explictly determining the distribution of θ from market data involving calibration quality is given in the following section.

As before, let $\mathbb{Q}$ denote the model that is used for hedging a European payoff X, and denote by $L = (L_t)_{0\le t\le T}$ the loss process when hedging with strategy Φ according to $\mathbb{Q}$. Losses from hedging in a misspecified model under model uncertainty have distribution function

$$\overline{\mathbb{P}}(L_t \le x) = \int_\Theta \overline{\mathbb{P}}_a(L_t \le x)\,\mu(\mathrm{d}a), \quad 0 \le t \le T. \tag{9}$$

This setup is compatible with the definition of a model-free strategy, see Definition 5. The notion of a model-free hedging strategy $\overline{\mathbb{P}}$–a.s., means that the respective strategy is model-free with respect to μ-almost all $\overline{\mathbb{P}}_a$, $a \in \Theta$.

Proposition 9. *A strategy Φ is a model-free hedging strategy for claim X $\overline{\mathbb{P}}$–a.s. if and only if $\overline{\mathbb{P}}(L_t = 0) = 1$, $0 \le t \le T$.*

Proof. "$\Rightarrow$": The claim follows directly from $\overline{\mathbb{P}}_a(L_t = 0) = 1$, for μ-a.a. $\overline{\mathbb{P}}_a$, $a \in \Theta$, and from $\overline{\mathbb{P}}(L_t = 0) = \mathbb{E}^{\overline{\mathbb{P}}}[\overline{\mathbb{P}}_\theta(L_t = 0)]$.

"$\Leftarrow$": Observe that $\overline{\mathbb{P}}_a(|L_t| > 0) \ge 0$, for μ-a.a. $\overline{\mathbb{P}}_a$, since $\overline{\mathbb{P}}_a$, $a \in \Theta$, are probability measures. Furthermore, $\overline{\mathbb{P}}(|L_t| > 0) = \mathbb{E}^{\overline{\mathbb{P}}}\left[\overline{\mathbb{P}}_\theta(|L_t| > 0)\right] = 0$ by assumption, so that $\overline{\mathbb{P}}_a(|L_t| > 0) = 0$ for μ-a.a. $\overline{\mathbb{P}}_a$ follows, see e.g. Section 6.2, Property F of Shiryaev (1996). □

Remark 10. Could one – as a consequence of Proposition 9 – use the measure $\overline{\mathbb{P}}$ as a model and derive an equivalent martingale measure for pricing and hedging to achieve model-free hedges? First, this would require knowledge of $\mathcal{P}$, resp. the set of equivalent martingale measures $\mathcal{Q}$, when setting up the hedging strategy and as such would be a violation of Assumption 3. As such, this would fail to capture the interpretation of $\mathcal{P}$ and $\mathcal{Q}$ as proxies for model uncertainty. Second – and a further justification of Assumption 3 –, this would be infeasible in practice. Consider the simple example of a market consisting of one asset and model set $\mathcal{Q} = \{\overline{\mathbb{Q}}_{\sigma_1}, \overline{\mathbb{Q}}_{\sigma_2}\}$ where the asset follows a Geometric Brownian motion with volatility σ_i in $\overline{\mathbb{Q}}_{\sigma_i}$, $i = 1, 2$. Provided that $\sigma_1 \ne \sigma_2$, the measures are singular, so that, at least theoretically, a claim can be perfectly hedged in each model, giving rise to a "model-free" hedging strategy. In practice, of course, such a strategy cannot be attained, since it requires determining the model based on the observed quadratic variation and then hedging in that model.

4.4 Model weights via AIC

A concrete approach to determine the distribution of θ is to use information about the calibration quality of each model in $\mathcal{P}$, resp. $\mathcal{Q}$. Such an approach uses only market information. Calibration quality requires calculating model prices as risk-neutral expectations, and it is therefore natural to work on the set $\mathcal{Q}$. For simplicity we shall assume in the following that there is a bijection between models $\overline{\mathbb{P}}_a$, $a \in \Theta$, and $\overline{\mathbb{Q}}_a$, $a \in \Theta$, so that model weights are equivalent for $\mathcal{P}$ and $\mathcal{Q}$.[7]

Essentially, a model with a smaller calibration error receives a greater probability weight than a model with a greater calibration error. However, it is always possible to improve the calibration quality by increasing the number of model parameters, which can easily lead to overly complex models, overfitting and poor robustness. So-called *model selection criteria* circumvent

[7] In a general setting, for each $\hat{\mathbb{P}} \in \mathcal{P}$, there may be one or more $\hat{\mathbb{Q}} \in \mathcal{Q}$, depending on whether $\hat{\mathbb{P}}$ determines a complete market or not. Conversely, for each $\hat{\mathbb{Q}} \in \mathcal{Q}$, there may be several $\hat{\mathbb{P}} \in \mathcal{P}$, e.g. for different market price of risk assumptions.

these problems by including, aside from a goodness-of-fit term such as calibration quality, a penalty term reflecting the model complexity.

A simple and popular model selection criterion is the *Akaike Information Criterion (AIC)* (Akaike, 1973). Under some regularity conditions, the AIC is an asymptotic unbiased estimate of expected relative discrepancy, where discrepancy is measured via Kullback-Leibler divergence, a measure of information loss between probability measures (i.e. models).[8] Akaike (1973) shows that the resulting estimator corresponds to the maximum likelihood of a model and a bias correction term, which originates from the lack of knowledge about which model constitutes the "true" one. The AIC is typically stated as a rescaled version of the above-mentioned estimator, given by

$$\mathrm{AIC} = -2\ln(\mathcal{L}) + 2K, \tag{10}$$

where $\mathcal{L}$ is the maximum of the likelihood function for the model and K is the number of parameters.[9] Given the data, a model with a smaller AIC has a smaller expected information loss and as such is preferred over a model with a higher AIC. For small sample sizes, a correction term applies, leading to AIC_c, given by $\mathrm{AIC}_c = \mathrm{AIC} + \dfrac{2K(K+1)}{n-K-1}$, where n is the sample size (Hurvich and Tsai, 1989). For more information on the AIC and model selection in general we refer to Gourieroux and Monfort (1995), Cavanaugh (1997) and Burnham and Anderson (2002, 2004).

The following Proposition shows how the mean square error (MSE), which is a commonly used measure of model calibration, see e.g. (Schoutens, 2003), can be related to AIC, thus linking calibration quality and the number of unknown parameters (as a bias correction term) in a unified criterion.

Proposition 11. *Let $C_1, \ldots, C_I$ be market prices of tradable benchmark options with payoffs $H_1, \ldots, H_I$, and let the mean-square error (MSE) of model $\overline{\mathbb{Q}}_a$ be given by*

$$MSE_a := \frac{1}{I}\sum_{i=1}^{I} |C_i - \mathbb{E}^{\overline{\mathbb{Q}}_a}[H_i]|^2. \tag{11}$$

Then MSE_a is the (quasi-)maximum likelihood of model $\overline{\mathbb{Q}}_a$ under the assumption that $\varepsilon_i := C_i - \mathbb{E}^{\overline{\mathbb{Q}}_a}[H_i]$, MSE_a, $i = 1, \ldots, I$, are independent, identically distributed realisations of a normal distribution with mean zero, and the corresponding AIC is given by

$$AIC(a) = I\,[1 + \ln(2\pi) + \ln(MSE_a)] + 2(K(a) + 1), \tag{12}$$

where $K(a)$ is the number of parameters in model $\overline{\mathbb{Q}}_a$.

The notion "quasi-maximum likelihood" refers to the fact that the assumption of normally distributed errors, although not justified, produces a maximum likelihood estimator (MLE) equivalent to minimising the MSE corresponding to a particular model family (e.g. Black-Scholes), and as such, the MLE has the same properties as the MSE.

Proof. To see that MSE_a is the maximum likelihood of model $\overline{\mathbb{Q}}_a$, it suffices to observe that MSE_a is the sample variance of the error terms $\varepsilon_1, \ldots, \varepsilon_I$ and corresponds to the MLE of the error terms' variance, given the parameters of model $\overline{\mathbb{Q}}_a$. The expression for the AIC is derived by simplifying the likelihood of model $\overline{\mathbb{Q}}_a$. The number of parameters entering the AIC corresponds to the number of model parameters and the variance parameter of the error terms. □

[8] The adjective "relative" relates to the fact that some terms of the estimator are constant across all models and are therefore dropped from the AIC, since they do not contribute to the model selection process.

[9] In the derivation of the penalty term in Equation (10) it is assumed that the "true" model belongs to the set according to which the MLE is determined. Takeuchi (1976) develops a model selection criterion, the *Takeuchi Information Criterion (TIC)*, that does not require this assumption. The AIC, however, is the criterion that is most widely used.

The number of model parameters $K(a)$ corresponds to the number of parameters entering the calibration procedure. For example, in a Black-Scholes model, the implied volatility enters as the single parameter, whereas for a Heston model five parameters are calibrated. A local volatility model can be uniquely calibrated only with an infinite number of options, so one needs to make additional assumptions, such as calibrating to piecewise constant integrated variance, in which case the number of parameters is just the number of options available.

If none of the models calibrates perfectly, then a probability distribution based on AIC is obtained via the likelihood of each model given the data, $\exp(-\Delta_a/2)$, where $\Delta_a = \mathrm{AIC}(a) - \mathrm{AIC}_{\min}$ and $\mathrm{AIC}_{\min} = \min_{\overline{\mathbb{Q}}_a \in \mathcal{Q}} \mathrm{AIC}_a$, cf. Burnham and Anderson (2004) (in case the minimum does not exist, one can use the infimum instead) and normalizing, to yield

$$\overline{\mathbb{P}}(\theta \in \mathrm{d}a) = \frac{\exp\{-\frac{1}{2}\Delta_a\}}{\int_{a\in\Theta} \exp\{-\frac{1}{2}\Delta_a\}\,\mathrm{d}\nu(a)}, \tag{13}$$

where ν is assumed to be a finite measure on Θ, such as the counting measure if $\mathcal{Q}$ is finite or Lebesgue measure if $\nu(\Theta) < \infty$.

When there are one or several models that calibrate perfectly, then the AIC, Equation (12), is $-\infty$ and Δ_a is ill-defined. Hence, this approach works only if there are sufficiently many benchmark options relative to the maximum number of parameters, and fails for non-parametric models. In practice, some seemingly non-parametric models such as local volatility models are often fitted to some analytic functional form, for example to achieve smoothness or to ensure absence of arbitrage, typically yielding a non-zero MSE (e.g., Brigo and Mercurio, 2002; Gatheral, 2006).

A criterion similar to AIC is the the *Bayesian Information Criterion (BIS)*, introduced by Schwarz (1978). BIC, given by $\mathrm{BIC} = -2\ln(\mathcal{L}) + K\ln(n)$, is similar to AIC, but with a larger penalty for the number of parameters, which is due to the assumption of uniform prior model weights. More generally, Bayesian model averaging methods are described in Hoeting, Madigan, Raftery, and Volinsky (1999).

In addition to market-related information one could use historical data to generate more refined probability weights. For example, using historical P&L from model risk would yield an improved discrimation of models based on their historical hedging performance, rather than relying on market price information alone.

5 Measures of model risk

We are now in a position to introduce measures of model uncertainty when pricing and hedging according to model $\mathbb{Q}$. As before, $Y := X - \sum_{i=1}^{I} u_i H_i$ and $\Phi = ((\phi(Y))_{0\le t\le T}, u_1, \dots, u_I)$ and

$$L_t(X,\Phi) = -\left(\mathbb{E}[Y] + \int_0^t \phi(Y)\,\mathrm{d}S - \mathbb{E}[Y|\mathcal{F}_t]\right) \tag{14}$$

denotes the time-t loss from hedging the claim X under $\mathbb{Q}$. Since $L_t(X,\Phi)$ is defined only $\mathbb{Q}$–a.s. and the measures in $\mathcal{P}$ are not necessarily absolutely continuous with respect to $\mathbb{Q}$, we shall always assume that concrete versions of prices and hedging strategies fulfilling Assumption 3 are chosen.

5.1 Value-at-risk and expected shortfall from model risk

The usual Value-at-risk and Expected Shortfall measures (e.g., McNeil, Frey, and Embrechts, 2005) are given as follows:

Definition 12. Let $L_t(X,\Phi)$ be the time-t loss from the strategy Φ that replicates claim X under $\mathbb{Q}$. Given a confidence level $\alpha \in (0,1)$,

(i) *Value-at-risk (VaR)* is given by $\mathrm{VaR}_\alpha(L_t(X,\Phi)) = \inf\{l \in \mathbb{R} : \overline{\mathbb{P}}(L_t(X,\Phi) > l) \le 1-\alpha\}$, that is, VaR_α is just the α-quantile of the loss distribution;

(ii) *Expected shortfall (ES)* is given by $\mathrm{ES}_\alpha(L_t(X,\Phi)) = 1/(1-\alpha)\int_\alpha^1 \mathrm{VaR}_u(L_t(X,\Phi))\,\mathrm{d}u$.

In the presence of benchmark instruments, the hedging strategy in model $\mathbb{Q}$ may not be unique. For every combination of static positions $u_1,\dots,u_I$ in the benchmark instruments, a version ϕ of the replicating strategy is chosen (cf. Assumption 3).

Let

$$\Pi(X) = \{\Phi : (u_1,\dots,u_I) \in \mathbb{R}^I \text{ and } \phi = \phi(Y) \text{ with } Y = X - \textstyle\sum_{i=1}^I u_i H_i\}$$

be the set of hedging strategies for claim X in model $\mathbb{Q}$. Here, the unique version of the strategy ϕ must fulfill Assumption 3.

To abstract from the particular hedging strategy chosen, we define measures that quantify the minimal degree of model dependence, indicating that when pricing and hedging under measure $\mathbb{Q}$, the model dependence cannot be further reduced. This is reasonable in the sense that it is not of interest whether a position is indeed hedged or not. Rather the hedging argument serves only to split the overall P&L into P&L from market risk and from model risk. In particular, choosing the measure that minimises model risk allows to appropriately capture claims that can be replicated in a model-free way.

The following defines measures of model risk similar to well-known risk measures for market risk or credit risk. However, since the full loss distribution is specified, any distribution-based risk measure may be defined.

Definition 13. Concrete measures capturing the model uncertainty when pricing and hedging claim X according to model $\mathbb{Q}$ are given by

(i) $\mu^{\mathbb{Q}}_{SQE,t}(X) = \inf_{\Phi\in\Pi(X)} \mathbb{E}[L_t(X,\Phi)^2]$,

(ii) $\mu^{\mathbb{Q}}_{\mathrm{VaR},\alpha,t}(X) = \inf_{\Phi\in\Pi(X)} \mathrm{VaR}_\alpha(|L_t(X,\Phi)|)$,

(iii) $\mu^{\mathbb{Q}}_{\mathrm{ES},\alpha,t}(X) = \inf_{\Phi\in\Pi(X)} \mathrm{ES}_\alpha(|L_t(X,\Phi)|)$.

(iv) $\rho^{\mathbb{Q}}_{\mathrm{VaR},\alpha,t}(X) = \inf_{\Phi\in\Pi(X)} \max(\mathrm{VaR}_\alpha(L_t(X,\Phi)),0)$,

(v) $\rho^{\mathbb{Q}}_{\mathrm{ES},\alpha,t}(X) = \inf_{\Phi\in\Pi(X)} \max(\mathrm{ES}_\alpha(L_t(X,\Phi)),0)$.

The measure $\mu^{\mathbb{Q}}_{SQE,t}$ is a simple measure of squared deviation of losses. The measures $\mu^{\mathbb{Q}}_{\mathrm{VaR},\alpha,t}$ and $\mu^{\mathbb{Q}}_{\mathrm{ES},\alpha,t}$ do not discriminate between profits and losses, but capture model uncertainty in an absolute sense. They can be thought of as as measures of the magnitude or degree of model uncertainty. The measures $\rho^{\mathbb{Q}}_{\mathrm{VaR},\alpha,t}$ and $\rho^{\mathbb{Q}}_{\mathrm{ES},\alpha,t}$ ignore potential profits and consider losses only. As such, they are suitable for defining a capital charge against losses from model risk.

Finally, to define measures of model uncertainty that depend solely on the claim but not on the particular measure used for pricing and hedging, one would first define the set $\overline{\mathcal{Q}} \subseteq \mathcal{Q}$ of potential pricing and hedging measures (e.g. measures that calibrate sufficiently well) and then define the risk measure in a worst-case sense as follows:

Definition 14. Let $\mu^{\hat{\mathbb{Q}}}_t(X)$ be a measure of model uncertainty when pricing and hedging X according to model $\hat{\mathbb{Q}} \in \overline{\mathcal{Q}}$. The model uncertainty of claim X is given by

$$\mu_t(X) = \sup_{\hat{\mathbb{Q}}\in\overline{\mathcal{Q}}} \mu^{\hat{\mathbb{Q}}}_t(X). \tag{15}$$

5.2 Capital charge for model risk

To provision against losses, either within an institution's risk policy or in terms of a regulatory capital charge, the measures $\rho^{\mathbb{Q}}_{\mathrm{VaR},\alpha,t}$ and $\rho^{\mathbb{Q}}_{\mathrm{ES},\alpha,t}$ serve as suitable candidates. First, when pricing and hedging in a model-dependent way, it makes sense to calculate provisions relative to the model used. Second, these risk measures are compatible with the respective risk measures for market risk used in practice, as they measure risk in the same risk units, in particular, if the time horizon t and confidence level α are the same. This curbs potential incentives to decrease market risk at the expense of increasing model risk.

One could argue that, instead of determining a capital charge based on the optimal hedging strategy, the capital charge should be calculated relative to the actual hedging strategy used. There are several reasons why this may not be practicable: First, when hedging a whole portfolio (e.g. the whole trading book), then it is not clear how to break down the overall hedging strategy to hedges for the individual positions due to diversification effects within the portfolio.[10] Second, although our approach depends crucially on the loss from hedging in a model-dependent way, model risk is present regardless of whether a position is hedged or not, and should be quantified as such. This justifies abstracting from the actual hedging strategy used, while not abstracting from the actual model used for pricing and hedging.

Of course, taking the infimum of associated losses can potentially incentivise to actively generate P&L from hedging in a more model-dependent way rather than a more defensive way, while this additional risk is not captured by e.g. risk limits or secured with capital. To rule out potential moral hazard issues, the actual amount of model risk can always be determined from a position's actual hedging strategy, or by enforcing a hedging policy.

The choice of the above risk measures explicitly rules out negative capital charges, although VaR and ES figures may be negative. One can choose a risk measure that explicitly allows for negative capital charges, reflecting for example that a model-dependent hedging strategy in model $\mathbb{Q}$ can act as a superhedge in all other models $\hat{\mathbb{Q}} \in \mathcal{Q}$.

5.3 Axioms for measures of model risk

We proceed to show that the measures introduced above fulfil some minimal desirable properties. Such axioms for measures of model uncertainty were formulated by Cont (2006). These axioms are based on the notion of coherent risk measures (Artzner, Delbaen, Eber, and Heath, 1999), and convex risk measures (Föllmer and Schied, 2002; Frittelli and Rosazza Gianin, 2002), which are widely accepted, but rather than postulating monotonicity and translation invariance, which make little sense in the context of model risk, the axioms for model risk measures take into account the possibility of hedging in a static way.

We adjust the axioms to account for potential losses from hedging realized prior to maturity of the option and for the fact that a particular model is chosen for hedging; see especially Axiom 2 below.

Cont (2006) assumes each model in the set $\mathcal{Q}$ expressing the model uncertainty to calibrate to the market in the sense that the prices of benchmark instruments are recovered within their bid-ask ranges. In our setting, this restriction is loosened, allowing explicitly for models that do not calibrate perfectly. However, a probability measure for the models should at the very least account for calibration quality.

Let further $\mathbb{Q}$ be a measure selected for pricing. Then, a mapping $\mu : \mathcal{C} \mapsto [0, \infty]$ representing model uncertainty should fulfill the following properties:[11]

1. For the liquidly traded benchmark instruments, model uncertainty is bounded by the bid-ask spread:
$$\forall i \in I, \quad \mu(H_i) \leq |C_i^{\text{ask}} - C_i^{\text{bid}}|. \tag{16}$$

[10]Of course, one could calculate the model risk associated with the overall portfolio and then apply techniques of capital allocation to break down the overall capital to individual positions, (e.g., Denault, 2001; Kalkbrener, 2005).

[11]Further we allow the measure to be infinite in order to express extreme degrees of model dependency. This turns out to be necessary to derive expected shortfall type measures of model risk.

2. Let $X \in \mathcal{C}$ a claim such that

$$\hat{\mathbb{Q}}\left(\mathbb{E}[X|\mathcal{F}_t] = \mathbb{E}[X|\mathcal{F}_0] + \int_0^t \phi_t(X)\,\mathrm{d}S_t\right) = 1, \forall \hat{\mathbb{Q}} \in \mathcal{Q}, \forall t, \tag{17}$$

   where $\mathbb{E}[X|\mathcal{F}_t]$ and $\phi_t(X)$ refers to the pricing functions and hedge positions chosen in Assumption 3. Then $\mu(X) = 0$. A claim that can be perfectly hedged across all models has no model risk. This in particular includes the case where X is defined in terms of a trading strategy. Further, let $\tilde{X} \in \mathcal{C}$ be another claim; then $\mu(X + \tilde{X}) = \mu(\tilde{X})$.
3. Convexity: diversification can decrease model uncertainty, that is,

$$\forall X, \tilde{X} \in \mathcal{C}, \forall \lambda \in [0,1] \quad \mu\left(\lambda X + (1-\lambda)\tilde{X}\right) \leq \lambda \mu(X) + (1-\lambda)\mu(\tilde{X}). \tag{18}$$

4. Static hedging with traded options decreases model uncertainty:

$$\forall X \in \mathcal{C}, \forall u \in \mathbb{R}^I, \quad \mu\Big(X + \sum_{i=1}^I u_i H_i\Big) \leq \mu(X) + \sum_{i=1}^I |u_i(C_i^{\text{ask}} - C_i^{\text{bid}})| \tag{19}$$

   In particular, if a payoff can be statically hedged with traded options, then the model uncertainty is bounded by the uncertainty on the cost of replication:

$$\left[\exists u \in \mathbb{R}^I, \quad X = \sum_{i=1}^I u_i H_i\right] \Rightarrow \mu(X) \leq \sum_{i=1}^I |u_i||C_i^{\text{ask}} - C_i^{\text{bid}}|. \tag{20}$$

Recall that in our setting, we ignore the upfront P&L from price discrepancies in the benchmark instruments. This P&L is realized immediately and as such treated as a sunk cost, so that a risk measure of model uncertainty captures only the uncertainty associated with future P&L. This allows to include models in the analysis that do not calibrate perfectly to the market, which is the de-facto standard even in practice. In fact, requiring models to calibrate perfectly is in conflict with the objective of model parsimony to prevent overfitting, which was discussed in Section 4.4. Since P&L from price discrepancies due to bid-ask spreads are up-front P&L, Axiom 1 reduces to $\mu(H_i) = 0$ and Axiom 4 reduces to $\mu\left(X + \sum_{i=1}^I u_i H_i\right) = \mu(X)$ since

$$\mu(X) = \mu\left(X + \sum_{i=1}^I u_i H_i - \sum_{i=1}^I u_i H_i\right) \leq \mu\left(X + \sum_{i=1}^I u_i H_i\right) \leq \mu(X). \tag{21}$$

The following Lemma captures the necessary ingredients for measures of model uncertainty when pricing and hedging according to model $\mathbb{Q}$. Let $\mathcal{L}$ be the linear space of functions $\Omega \to \mathbb{R}$.

Lemma 15. *Let $f : \mathcal{L} \to [0,\infty]$ be a convex function satisfying $f(L) = 0$ whenever $L = 0$ $\overline{\mathbb{P}}$–a.s. and $f(L) = f(\tilde{L})$ whenever $L = \tilde{L}$ $\overline{\mathbb{P}}$–a.s.. Then $\mu_t^{\mathbb{Q}} : \mathcal{C} \to [0,\infty]$ given by*

$$\mu_t^{\mathbb{Q}}(X) = \inf_{\Phi \in \Pi} f(L_t(X, \Phi)) \tag{22}$$

is a measure satisfying the axioms of model uncertainty. If f is not convex, then $\mu_t^{\mathbb{Q}}(X)$ satisfies Axioms 1, 2 and 4.

The proof is in Appendix C. It should be noted here that Assumption 3 is assumed to hold.

Proposition 16. *The measures $\mu^{\mathbb{Q}}_{SQE,t}(X), \mu^{\mathbb{Q}}_{ES,\alpha,t}(X)$ and $\rho^{\mathbb{Q}}_{ES,\alpha,t}(X)$ satisfy the axioms of model uncertainty. The measures $\mu^{\mathbb{Q}}_{VaR,\alpha,t}(X)$ and $\rho^{\mathbb{Q}}_{VaR,\alpha,t}(X)$ satisfy Axioms 1, 2 and 4.*

Proof. For each measure it is easily seen that the respective function defining the measure fulfills $f(L) = 0$ whenever $L = 0$ $\overline{\mathbb{P}}$–a.s. and $f(L) = f(\tilde{L})$ whenever $L = \tilde{L}$ $\overline{\mathbb{P}}$–a.s.. For the square error and expected shortfall based measures, we must in addition show that the respective functions are convex. For (i) of Definition 13, it is easily shown that $\mathbb{E}[\xi^2]$ is convex for any random variable ξ with finite second moment. For (iii), it suffices to observe that expected shortfall is convex (see Proposition 6.9 of McNeil et al. (2005) for a proof), and for (v) in addition it is easily shown that $\max(g(x), 0)$ is convex if $g(x)$ is convex. Furthermore, by using linearity of expectation it can easily be verified that convexity extends to the region where the respective functions attain ∞. □

Proposition 17. *The measure $\mu_t(X)$ fulfills Axioms 1, 2 and 4. If $\mu_t^{\hat{\mathbb{Q}}}(X)$ fulfills Axiom 3 for all $\hat{\mathbb{Q}} \in \overline{\mathcal{Q}}$, then $\mu_t(X)$ fulfills Axiom 3.*

Proof. Axioms 1 and 2 hold trivially as they hold for all $\hat{\mathbb{Q}} \in \mathcal{Q}$.

For Axiom 3, suppose that $\mu_t^{\hat{\mathbb{Q}}}(X)$ fulfills Axiom 3 for all $\hat{\mathbb{Q}} \in \overline{\mathcal{Q}}$, and let $X, \tilde{X} \in \mathcal{C}$ and $\lambda \in [0, 1]$. Then,

$$\begin{aligned}
\mu_t(\lambda X + (1-\lambda)\tilde{X}) &= \sup_{\hat{\mathbb{Q}} \in \overline{\mathcal{Q}}} \mu_t^{\hat{\mathbb{Q}}}(\lambda X + (1-\lambda)\tilde{X}) \\
&\leq \sup_{\hat{\mathbb{Q}} \in \overline{\mathcal{Q}}} \left\{ \lambda \mu_t^{\hat{\mathbb{Q}}}(X) + (1-\lambda)\mu_t^{\hat{\mathbb{Q}}}(\tilde{X}) \right\} \\
&\leq \lambda \mu_t(X) + (1-\lambda)\mu_t(\tilde{X}),
\end{aligned} \tag{23}$$

where the first inequality follows from the convexity of $\mu_t^{\hat{\mathbb{Q}}}$ and the second inequality follows from properties of the supremum.

For Axiom 4, we have

$$\mu_t\left(X + \textstyle\sum_{i=1}^{I} u_i H_i\right) = \sup_{\hat{\mathbb{Q}} \in \overline{\mathcal{Q}}} \mu_t^{\hat{\mathbb{Q}}}\left(X + \textstyle\sum_{i=1}^{I} u_i H_i\right) \leq \sup_{\hat{\mathbb{Q}} \in \overline{\mathcal{Q}}} \mu_t^{\hat{\mathbb{Q}}}(X) = \mu_t(X). \tag{24}$$

□

6 Case studies

To investigate the magnitude of model risk, we calculate value-at-risk and expected shortfall in various settings.

6.1 Model risk under dynamic hedging

The first example considers a setting similar to the empirical example in Section 2. The goal is to determine the model risk of an at-the-money call option maturing in three months. There are no benchmark instruments in the market and the only hedging strategy is fully dynamic. The model set contains Black-Scholes models of various implied volatilities. The volatility used for pricing and hedging is $\sigma = 25.401\%$, which corresponds to the average VIX volatility from the data set used in Section 2. Likewise, the risk-free interest rate is calculated as an average of $r = 2.064\%$. Despite the lack of benchmark instruments, model weights are calculated from calibration error relative to the model used for pricing and hedging.

We calculate model risk under various parameter constellations. In any of the examples below, the loss distribution is estimated with Monte Carlo simulation as follows: Observe first that the loss relative to any model is given by Equation (7). In each of $10,000$ simulations, the stock price path is simulated at an hourly frequency and the integral (7) is approximated accordingly. The loss distribution is given by a smooth kernel estimate of the simulated losses.

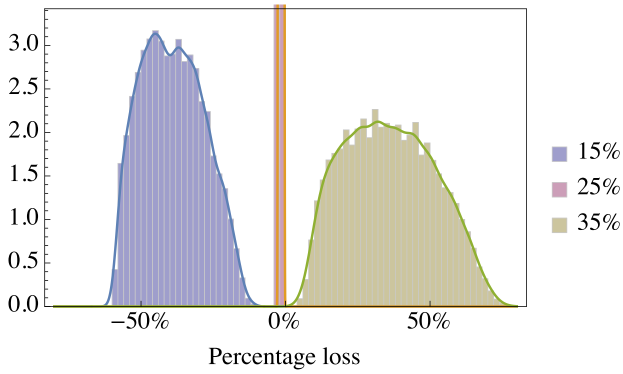

Figure 2: Loss distributions with respect to different model assumptions. The pricing model has a volatility of $\sigma = 0.254$. When the market follows a Black-Scholes model with volatility of 15% (blue), then hedging in the pricing model leads to a negative loss (profit), whereas when the market follows a model with volatility of 35% (green), there is a loss.

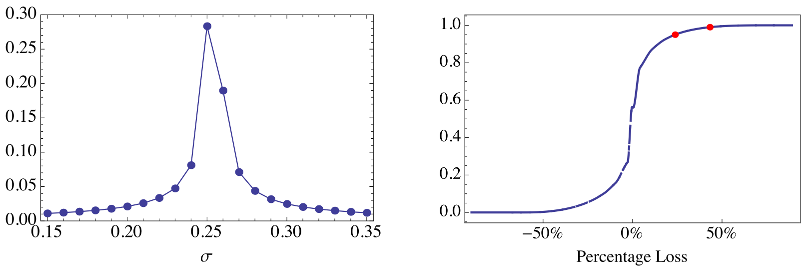

Figure 3: Left: Probability weights of models in set $\{15\%, 16\%, \dots, 35\%\}$ according to the AIC. Right: Final loss distribution, with 95%- and 99%-VaR marked by red points.

Model weights are calculated according to the AIC criterion (Section 4.4 and Equation (13)). The 99%- and 95%-VaRs are calculated from the unified loss distribution.

In the first examples, we fix the model set as the range of Black-Scholes models with implied volatilities in $\{15\%, 16\%, \dots, 35\%\}$. First, fix the drift μ at 5%. Figure 2 shows the loss densities in the cases where the market follows the extreme cases of 15% and 35% volatilities, and where the market volatility of 25% is near the volatility of the pricing model. The probability weights assigned to the models are shown in the left graph of Figure 3. The resulting overall loss distribution is shown in the right graph of Figure 3, where the red points mark the 95%- and the 99%-VaR, respectively.

Figure 4 shows value-at-risk at 99% and 95% levels as functions of the drift parameter μ and of the market price of risk $\lambda = (\mu - r)/\sigma$, which are kept fixed across the set of models. It turns out that the VaR numbers are insensitive to varying drift or market price of risk, so that the type of measure (objective or risk-neutral) is of minor importance.

In the next examples, the drift rate is fixed at $\mu = 0.05$, but the set of models quantifying the uncertainty varies. The models are chosen evenly-spaced in 1%-intervals around a volatility of 25%. The range of models is determined via a distance parameter to the 25% model. Figure 4 shows the resulting model risk figures as a function of distance. Not surprisingly, in the setup without benchmark instruments, model risk depends strongly on the set denoting the model uncertainty. However, for a more mature market with a higher number of benchmark instruments, one would expect the calibration quality to discriminate more strongly between the models thus leading to less dependence on the model set specification.

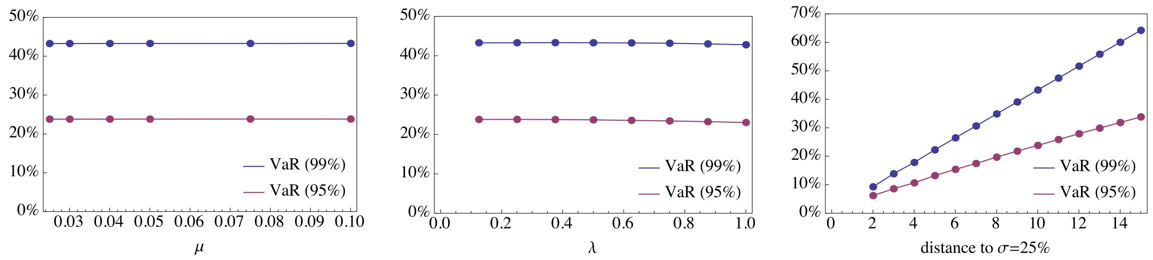

Figure 4: Left and middle: Value-at-risk for model risk from different values of drift (left) and market price of risk (middle). Model uncertainty is captured by a set of Black-Scholes models with volatilities ranging in 0.15 and 0.35. Right: Value-at-risk for model risk under different sets of Black-Scholes models quantifying the model uncertainty. The pricing model has an implied volatility of 0.254.

6.2 Model risk when including static hedging

In a more sophisticated example we include benchmark instruments implying the possibility of static hedging and we consider a more realistic set of models consisting of Black-Scholes model and Heston models with varying parameters. Assuming zero interest rates the dynamics of the Black-Scholes model are described by

$$\mathrm{d}S_t = \sigma S_t \mathrm{d}W_t,$$

where $W = (W_t)_{t \geq 0}$ is a Brownian motion and with constant volatility σ. The dynamics of the Heston model (Heston, 1993) are

$$\begin{aligned} \mathrm{d}S_t &= \sqrt{V_t} S_t \mathrm{d}W_{t,1} \\ \mathrm{d}V_t &= \kappa(\zeta - V_t)\mathrm{d}t + \nu\sqrt{V_t}\mathrm{d}W_{t,2}, \end{aligned}$$

with mean reversion level ζ, mean reversion rate κ and volatility of volatility ν. The instantaneous correlation of the two Brownian motions W_1 and W_2 is denoted by ρ.

All calculations are based on implied volatilities of options on the S&P 500 as published on Bloomberg on 15 May 2013, comprising 11 different strikes ranging from 80% to 120% of spot and maturities ranging from one month to two years. To calibrate the model parameters of each model to prices we minimize the root-mean-square deviation between model prices and market prices. Both the risk horizon and the maturity of the example payoffs considered is one year, so that the models are calibrated to the options with one maturity only. Since the mean reversion rate κ of the Heston model cannot be uniquely identified by options with the same maturity, the estimate for κ is based on all option prices and then enters the calibration restricted to one-year options. Table 1 shows the parameter estimates for both models. Denote by $\overline{\mathbb{Q}}_\sigma$ the Black-Scholes model with volatility parameter σ and by $\overline{\mathbb{Q}}_{V_0,\zeta,\kappa,\nu,\rho}$ the Heston model with its five parameters V_0, ζ, κ, ν and ρ. We build an example model set by

$$\begin{aligned} \mathcal{Q} \;:=\; & \{\overline{\mathbb{Q}}_\sigma | \sigma > 0\} \cup \{\overline{\mathbb{Q}}_{V_0,\hat{\zeta},\hat{\kappa},\hat{\nu},\hat{\rho}} | V_0 \in [0.016, 0.0175]\} \\ \cup \; & \{\overline{\mathbb{Q}}_{\hat{V}_0,\zeta,\hat{\kappa},\hat{\nu},\hat{\rho}} | \zeta \in [0.049, 0.051]\} \cup \{\overline{\mathbb{Q}}_{\hat{V}_0,\hat{\zeta},\hat{\kappa},\nu,\hat{\rho}} | \nu \in [0.545, 0.575]\} \end{aligned}$$

The model set is discretized with 45 equally spaced parameter values for each parameter type in the above domain. We determine the distribution on the model set according to the Akaike Information Criterion as described in Section 4.4. The set of Black-Scholes models is assigned a probability of zero due to its high calibration error relative to the Heston model. It turns out that the model density for the Heston models is concentrated on a very small domain

Model	Parameters				
Black-Scholes	$\hat{\sigma}$				
	0.1613				
Heston	$\hat{V}_0$	$\hat{\zeta}$	$\hat{\kappa}$	$\hat{\nu}$	$\hat{\rho}$
	0.0167	0.0501	1.6052	0.5600	-0.6243

Table 1: Parameters estimates for Heston and Black-Scholes models based on market prices of options on the S&P 500 from 15 May 2013 (Source: Bloomberg). Except for $\hat{\kappa}$ all estimates are based on options with a maturity of one year. The estimate $\hat{\kappa}$ is based on options with different maturities.

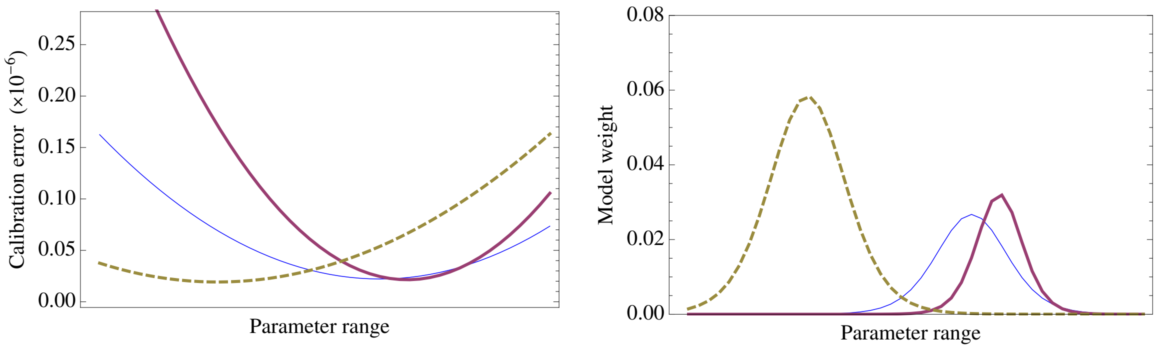

Figure 5: Left: Calibration error of the Heston models given by the root-mean-squared deviation. Right: AIC model weights for the Heston models as given by Equation (13). Heston models with parameters based on the calibration (see Table 1) except for one parameter that is changed: $V_0 \in [0.016, 0.0175]$ (thin solid line), $\zeta \in [0.049, 0.051]$ (thick solid line) and $\nu \in [0.545, 0.575]$ (dashed line).

and the parameter ranges in the definition of $\mathcal{Q}$ are chosen such that the complementary range has negligible probability. Adding models with parameters outside the parameter range used in the definition of $\mathcal{Q}$ does not change the results. This backs the conclusion from the last case study that more benchmark instruments are needed to sufficiently discriminate between the available models and to make the risk figures robust against a change of the model set. The left graph in Figure 5 shows the calibration error as measured by the root-mean-square deviation of the Heston models with varying parameters V_0, ϑ and ν. The graph on the right hand side shows the model probabilities given by the Akaike weights, Equation (13). For simplicity we assume $\mathcal{Q} = \mathcal{P}$, that is, we estimate the loss distribution in a risk-neutral setting.

For static hedging two liquidly traded benchmark options H_1 and H_2 are available, both of which have a maturity of one year. The option H_1 is a put option with strike $0.8S_0$ ($S_0 = 1$). H_2 is an at-the-money call option with strike 1. The corresponding observed market prices are $C_1 = 0.015$ and $C_2 = 0.064$, based on implied volatilities from the market data of 21.60% and 16.07%.

As a first example we measure the model risk of a short position in a one-year call option with strike 1.1, that is, with payoff $X = -(S_1 - 1.1)^+$. The pricing model $\mathbb{Q}$ corresponds to the Black-Scholes model with $\sigma = 0.142$, which is just the volatility implied by the market data, yielding a price of 0.022. First, we consider the four strategies $\Phi_1 = ((\phi_t^1)_{0\leq s\leq T}, 0, 0)$, $\Phi_2 = ((\phi_t^2)_{0\leq s\leq T}, 1, 0)$, $\Phi_3 = ((\phi_t^3)_{0\leq s\leq T}, 0, 1)$ and $\Phi_4 = ((\phi_t^4)_{0\leq s\leq T}, u_1, u_2)$, where the ϕ^i are such that $\mathbb{Q}(L_T(X, \Phi_i) = 0) = 1$ for $i \in \{1, 2, 3, 4\}$ and u_1, u_2 are the positions in the benchmark instruments that minimize the 95%-VaR $\rho^{\mathbb{Q}}_{\mathrm{VaR},0.95,1}$. Figure 6 shows Box-Whisker plots of the

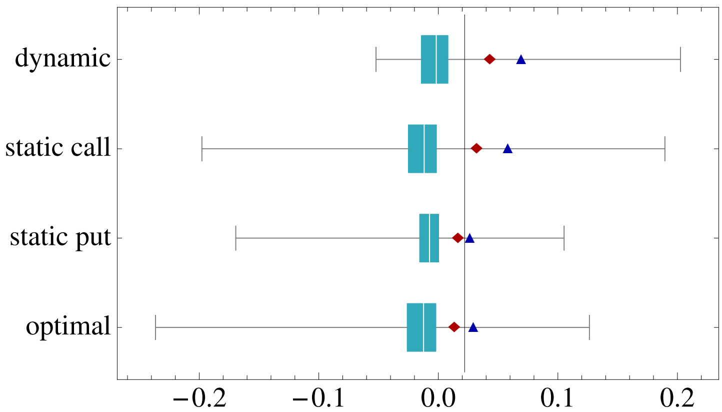

Figure 6: Box-whisker plots of loss distributions from hedging a short position in a call option with strike 1.1 under various hedging strategies. Each box and each whisker comprises 25% of the distribution, the diamonds denote the 95%-VaR, the triangles denote the 95%-ES and the vertical line corresponds to the option premium. Top: Purely dynamic hedging strategy; VaR is 192.22% and ES is 309.53% of option premium. Second: Hedging strategy involving static position in benchmark put option; VaR is 143.31% and ES is 259.93% of option premium. Third: Hedging strategy involving static position in benchmark call option; VaR is 73.39% and ES is 117.64% of option premium. Bottom: $\rho^{\mathbb{Q}}_{\mathrm{VaR},0.95,1}$-optimized strategy; VaR is 59.81% and ES is 130.46% of option premium.

distributions of $L_T(X,\Phi_1)$, $L_T(X,\Phi_2)$, $L_T(X,\Phi_3)$ and $L_T(X,\Phi_4)$ under $\overline{\mathbb{P}}$ based on $10,000$ simulations with $1,000$ time steps each. Even at a confidence level of 0.95, value-at-risk and expected shortfall turn out to be rather high for strategy Φ_1 with 0.0430 and 0.0692, which corresponds to 192.22% and 309.53% of the option value. The high risk figures are due to the fact that the pricing volatility is relatively low for the out-of-the-money call option because of the strong smile in the volatility. In addition, the long maturity of the option contributes to a greater possible loss under model misspecification.

The value-at-risk figures for strategies Φ_2 and Φ_3, given by 0.0321 and 0.0164, are significantly lower than for Φ_1. This can be attributed to the static hedge positions. In particular, the position $1H_2$ is close to the initial vega hedge, which is $0.84H_2$ (the vega hedge at time 0 with H_1 would be $3.21H_1$). The minimum VaR is achieved with positions of $1.15\,H_1$ and $0.63\,H_2$ yielding a VaR of 0.0133 and an ES of 0.02918.

Aside from the minimal VaR, Table 2 shows in addition the other risk measures from Definition 13. Since all risk measures are defined as the infimum over the feasible strategies for a claim, the static positions in the benchmark instruments H_1 and H_2 at the infimum are shown as well. Of course, depending on the risk measure chosen the risk minimizing strategy changes.

The second example is a short position in a digital put option paying 1 if the final spot value is below the strike of $K = 0.8$ and nothing otherwise, that is, with payoff $\tilde{X} = -\mathbf{1}_{\{S_T\leq 0.8\}}$. The pricing model $\mathbb{Q}$ corresponds to the Black-Scholes model with $\sigma = 0.216$, which is just the volatility implied by the market data for strike 0.8, yielding a price of 0.178. As before we consider the four strategies $\tilde{\Phi}_1 = ((\tilde{\phi}^1_t)_{0\leq s\leq T}, 0, 0)$, $\tilde{\Phi}_2 = ((\tilde{\phi}^2_t)_{0\leq s\leq T}, 1, 0)$, $\tilde{\Phi}_3 = ((\tilde{\phi}^3_t)_{0\leq s\leq T}, 0, 1)$ and $\tilde{\Phi}_4 = ((\tilde{\phi}^4_t)_{0\leq s\leq T}, \tilde{u}_1, \tilde{u}_2)$ with $\tilde{\phi}^i$ such that $\mathbb{Q}(L_T(\tilde{X},\tilde{\Phi}_i) = 0) = 1, i \in \{1,2,3,4\}$. Figure 7 shows Box-Whisker plots of the distributions of $L_T(\tilde{X},\tilde{\Phi}_i)$ under $\overline{\mathbb{P}}$. The risk figures are higher than in the last example with a value-at-risk of 0.174 and an expected shortfall of 0.389 for strategy $\tilde{\Phi}_1$. The benefit of partial hedging is weaker in this example. The optimal hedging strategies given in Table 2 indicate that the strategies $\tilde{\Phi}_2$ and $\tilde{\Phi}_3$ are not optimal in reducing value-at-risk and Expected shortfall. Depending on the risk measure, the optimal position u_1 in H_1 ranges from -0.57 to 8.53 and in H_2 from -0.58 to 3.71. The total holding $u_1 + u_2$ is

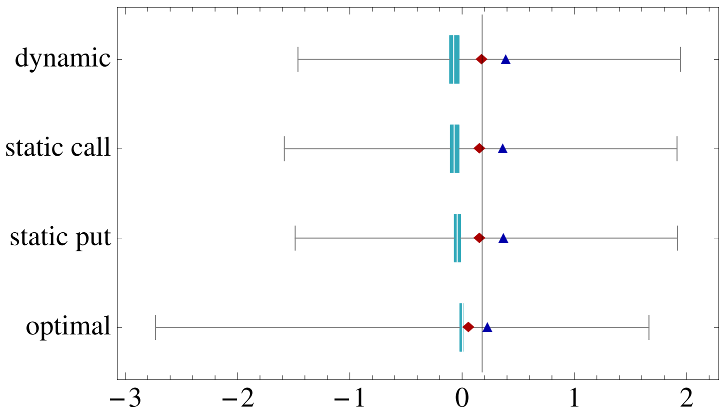

Figure 7: Box-whisker plots of loss distributions from hedging a short position in a digital option with strike 0.8 under various hedging strategies. Each box and each whisker comprises 25% of the distribution, the diamonds denote the 95%-VaR, the triangles denote the 95%-ES and the vertical line corresponds to the option premium. Top: Purely dynamic hedging strategy; VaR is 98.13% and ES is 219.13% of option premium. Second: Hedging strategy involving static position in benchmark put option; VaR is 87.07% and ES is 203.73% of option premium. Third: Hedging strategy involving static position in benchmark call option; VaR is 86.57% and ES is 207.23% of option premium. Bottom: $\rho^{\mathbb{Q}}_{\mathrm{VaR},0.95,1}$-optimized strategy; VaR is 31.48% and ES is 127.26% of option premium.

always greater than 2.98, which suggests that both strategies $\tilde{\Phi}_2$ and $\tilde{\Phi}_3$ have too small holdings in the benchmark options to reduce model risk. A vega hedge with H_1 would be a holding of $6.60H_1$, while with H_2 it would be $3.46H_2$. The risk measures are considerably higher than for the call option, suggesting a higher model dependence also of the optimal strategy due to the discontinuity of the payoff and less similarity to the benchmark options. Although the strike of the put option H_1 used for the static hedge is the same as the strike of $\tilde{X}$, the part that is dynamically hedged, $\tilde{X} - u_1H_1 - u_2H_2$, faces a greater model dependence. Both examples indicate that the market practice of using a Black-Scholes model with its implied volatility for simple options entails significant model risk.

6.3 Gap risk

As a final example, we consider the measurement of gap risk, a risk that is introduced by the presence of jumps in asset prices. In addition, the example demonstrates how a semi-static hedging strategy involving a finite number of trades at stopping times in the benchmark instruments can be incorporated in the model risk framework.

The setup is as follows: We consider a call option with a down-and-out barrier (DOC option) on the forward price of the underlying asset. The payoff is given by

$$(S_T - K)^+ \, \mathbf{1}_{\{\inf_{0\le t\le T} F_{t,T} > B\}},$$

where $S = (S_t)_{t\ge 0}$ is the price process of the underlying asset, $F_{t,T} = \mathbb{E}^{\mathbb{Q}}[S_T|\mathcal{F}_t]$ is the price process of the corresponding time-T forward, T is the maturity of the option, K its strike and B the barrier.

In the special case where the barrier and the strike coincide, $B = K$, and assuming a continuous price process, the following semi-static strategy hedges a short position in the DOC option (Carr, Ellis, and Gupta, 1998; Albrecher and Mayer, 2010): Buy a call option and sell a put option of the same maturity, both with a strike of $K = B$, and unwind the hedge portfolio when the barrier is hit. Because of put-call parity the value of the hedge portfolio will be zero

	$X=-(S_T-1.1)^+$				$\tilde{X}=-\mathbf{1}_{\{S_T\leq 0.8\}}$			
Measure	Absolute	Percentage	u_1	u_2	Absolute	Percentage	$\tilde{u}_1$	$\tilde{u}_2$
$\mu^{\mathbb{Q}}_{SQE,1}$	0.0002	0.92%	−0.15	0.71	0.0173	9.74%	0.29	2.49
$\mu^{\mathbb{Q}}_{\mathrm{VaR},\alpha,1}$	0.0275	122.73%	−0.18	0.74	0.0255	144.01%	−0.57	3.71
$\mu^{\mathbb{Q}}_{\mathrm{ES},\alpha,1}$	0.0354	158.27%	−0.23	0.73	0.4020	226.60%	−0.49	3.68
$\rho^{\mathbb{Q}}_{\mathrm{VaR},\alpha,1}$	0.0133	59.81%	1.15	0.63	0.0559	31.48%	8.53	0.70
$\rho^{\mathbb{Q}}_{\mathrm{ES},\alpha,1}$	0.0227	101.42%	0.95	1.31	0.1710	96.19%	13.45	−0.58

Table 2: Model risk as defined in Definition 13 with confidence level $\alpha = 0.95$ and time horizon $T = 1$. Claim X is a short position in a one year call option with strike 1.1 and claim $\tilde{X}$ a short position in a one year digital put option with strike 0.8. The risks are shown both as absolute figures and as a percentage of the initial option premium. u_i, resp. $\tilde{u}_i$, $i = 1, 2$, denotes the static holdings in benchmark instruments.

Parameter	value
Asset process and DOC option	
Stock price S_0	1
Strike K	0.75
Barrier B	0.75
Maturity T	5 years
Risk-free interest rate r	0%
Jump-diffusion parameters	
Diffusion drift μ	5%
Diffusion volatility $\hat{\sigma}$	0.10293
Jump rate $\hat{\lambda}$	0.23203
Jump mean $\hat{a}$	−0.18640
Jump volatility $\hat{b}$	0.28710
Simulation	
Number of simulations	10,000

Table 3: Parameters for estimating gap risk of DOC barrier option. The Merton jump-diffusion model σ, a, b and λ are derived from calibrating against market prices of options on the S&P 500 from 15 May 2013 (Source: Bloomberg).

when $F_{t,T} = B$. If the barrier is not hit during the lifetime of the DOC option, then both the DOC and the call option will be in the money at maturity, while the put option expires worthless.

In the presence of jumps, however, the strategy is no longer a hedge: in case of a jump beyond the barrier, the portfolio can be unwound only at a loss of $F_{t,T} - B$. This type of risk is called *gap risk*. We calculate this risk in the model risk framework.

First, we assume that the pricing model $\mathbb{Q}$ corresponds to a diffusion model, so that the semi-static strategy is indeed a hedge. Because of the model-independent hedging strategy, the precise specification of the model is irrelevant. In order to incorporate the semi-static hedge, we can either introduce the call and put option portfolio from the hedge as a separate asset price process, or we can specify the entire hedge portfolio strategy (call and put being unwound the first time $F_{t,T}$ hits B) as a benchmark instrument. Either choice leads to the same result.

Next, the set $\mathcal{P}$ is specified to contain models incorporating jumps in the asset price process, and the hedge error relative to those models is calculated. In our setup, we choose $\mathcal{P}$ to contain jump-diffusion models (Merton, 1976) with varying parameters around the calibrated model.

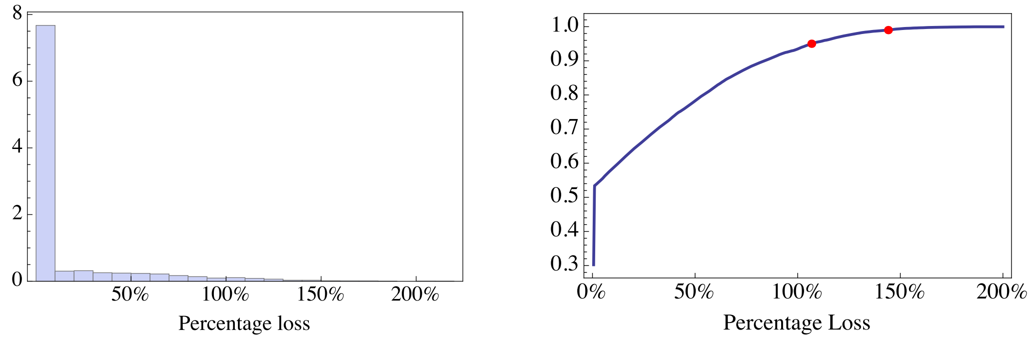

Figure 8: Loss distributions of gap risk in DOC barrier option. Loss (as percentage of option premium in diffusion model) from semi-static hedging strategy when asset price process follows jump-diffusion model. Left: Histogram of loss distribution in jump-diffusion model calibrated to market. Right: Aggregated cumulative loss distribution from models in $\mathcal{Q}$, with 95% and 99%-VaR marked by red points.

The jump-diffusion model postulates that asset prices are given by (under $\mathbb{P}$)

$$S_t = S_0 \exp\Big(\mu t + \sigma W_t + \sum_{i=1}^{N_t} Y_i\Big), \quad t \geq 0,$$

where W is a Brownian motion, N is a Poisson process with intensity λ, independent of W, and $Y_1, Y_2, \ldots$ are independent normally distributed random variables with mean a and standard deviation b, independent of W and N.

We calibrate the risk-neutral parameters σ, λ, a and b to prices of traded options on the S&P 500 on 15 May 2013, and assume that the real-world parameters are given by the risk-neutral parameters and the drift μ (that the measures are equivalent is established in Proposition 9.8 of Cont and Tankov (2004), resp. Theorems 33.1/33.2 of Sato (1999)).

Table 3 contains the parameters of the market and DOC barrier option as well as the parameters of the calibrated jump-diffusion model. Because of the assumption of a zero risk-free interest rate, the asset price and the forward price process are equivalent. The model set $\mathcal{P}$ contains the models

$$\begin{aligned}\mathcal{P} := &\{\overline{\mathbb{P}}_{\sigma,\hat{\lambda},\hat{a},\hat{b}} | \sigma - \hat{\sigma} \in \{-0.03, -0.027, -0.024, \ldots, 0.027, 0.03\}\} \\ &\cup \{\overline{\mathbb{P}}_{\hat{\sigma},\lambda,\hat{a},\hat{b}} | \lambda - \hat{\lambda} \in \{-0.1, -0.09, -0.08, \ldots, 0.09, 0.1\}\} \\ &\cup \{\overline{\mathbb{P}}_{\hat{\sigma},\hat{\lambda},a,\hat{b}} | a - \hat{a} \in \{-0.03, -0.027, -0.024, \ldots, 0.027, 0.03\}\} \\ &\cup \{\overline{\mathbb{P}}_{\hat{\sigma},\hat{\lambda},\hat{a},b} | b - \hat{b} \in \{-0.1, -0.09, -0.08, \ldots, 0.09, 0.1\}\}\end{aligned}$$

The loss distribution from the semi-static hedging strategy when the asset price process follows a jump-diffusion is determined by Monte Carlo simulation. Fixing a jump-diffusion model, in each simulation scenario, first the jump times $\tau_1, \tau_2, \ldots$ are determined. In addition, the minimum of the stock price between any two jump times τ_{i-1}, τ_i, conditional on the stock prices $S_{\tau_{i-1}}$, $S_{\tau_i -}$ is simulated. When considering logs of stock prices, this corresponds to simulating the minimum of a Brownian bridge with drift, given the distribution function (Borodin and Salminen, 2002, 1.2.8 on p. 252) and via the inverse transform method (Glasserman, 2004, Section 2.2.1). Whenever the barrier is first hit by a jump, P&L is generated, otherwise it is checked whether the barrier is hit by the diffusion part in-between two jumps, in which case no P&L is generated. The loss distribution for the calibrated jump-diffusion model (that is, with parameters given by Table 3) is shown in the left graph of Figure 8. The expected loss is 13.75% of the option premium; 95%-VaR is 86% and 99%-VaR is 128% of the option premium.

The cumulated distribution function of the AIC-weighted aggregated loss distribution from the set $\mathcal{P}$ is shown in the right graph of Figure 8. Here, the 95%-VaR is 107% and the 99%-VaR is 144%.

7 Conclusion and outlook

An appropriate assessment of model risk when trading contingent claims is important for several reasons: First, assessing the potential losses associated with a claim from using a model for pricing adds to the proper understanding of risks in trading books beyond market risk. Second, revealing potentially high losses from model uncertainty inherent in a position can prevent unintentional risk-taking and associated risk-taking-related incentive conflicts. Third, an adequate assessment of model risk is suitable for deriving capital requirements against unexpected losses from model risk.

Model uncertainty is expressed via a set of models $\mathcal{P}$ all of which determine suitable models for the asset price processes. Associated with $\mathcal{P}$ is a set $\mathcal{Q}$ of equivalent martingale measures, giving rise to suitable pricing measures. Given the model used for pricing, a claim's potential losses from model risk are captured via the "residual P&L" assuming that the claim is perfectly hedged, which in a complete market is equivalent to eliminating the claim's market risk. In a first step, we derive an expression for this loss relative to only one model, sometimes called tracking error in the existing literature.

We equip the set $\mathcal{P}$ with a probability distribution, so that the probability measures in $\mathcal{P}$ form a regular conditional probability on an extended probability space. This allows to derive an aggregated loss distribution for the losses from model risk taking into account all models in $\mathcal{P}$. The Akaike Information Criterion (AIC) provides one method of deriving a probability distribution of the models in $\mathcal{P}$, resp. $\mathcal{Q}$. Only market information, via the calibration error and model complexity enter the AIC-derived probabilities. One could further incorporate historical data, e.g. on the hedge quality, to refine the probability weights.

Given the loss distribution from model uncertainty, value-at-risk (VaR) and expected shortfall (ES) measures for model risk are defined in the usual way. In case of several hedging strategies due to the possibility of static hedging with liquidly traded options, the smallest VaR or ES is chosen to quantify model risk. This allows to extend the notion of model risk to unhedged positions. The measures proposed fulfil the axioms for measures of model uncertainty formulated by Cont (2006) (with the usual exception of subadditivity for VaR). Static hedging with liquidly traded options reduces model uncertainty and as such should be preferred over dynamic, model-dependent hedging strategies.

Several case studies demonstrate the magnitude of model risk. We consider the setup where only dynamic hedging is possible and determine the sensitivity of model risk to the drift, resp. market price of risk, and to the model set $\overline{\mathbb{P}}$. Next, we investigate the change in model risk between dynamic hedging and static hedging. Finally, we consider semi-static hedging strategies and demonstrate a method for calculating gap risk. In all examples, we find that model risk is significant and can even exceed the option premium.

A The price range measure

The price range measure is a simple and popular indicator of model risk. Across a set of pricing models $\mathcal{Q}$, the price range of a claim X, which is assumed to have a well-defined price in each model, is given by (Cont, 2006):

$$\mu_{\mathcal{Q}}(X) = \sup_{\hat{\mathbb{Q}}\in\mathcal{Q}} \mathbb{E}^{\hat{\mathbb{Q}}}[X] - \inf_{\hat{\mathbb{Q}}\in\mathcal{Q}} \mathbb{E}^{\hat{\mathbb{Q}}}[X]. \tag{25}$$

Certainly, if all measures calibrate perfectly to the same benchmark instruments, then a payoff whose value is not influenced by model uncertainty, for example because of the possibility of static hedging, has $\mu_{\mathcal{Q}}(X) = 0$. A margin or provision for model uncertainty when using model $\mathbb{Q}$ for pricing is given by $\sup_{\hat{\mathbb{Q}}\in\mathcal{Q}} \mathbb{E}^{\hat{\mathbb{Q}}}[X] - \mathbb{E}^{\mathbb{Q}}[X]$.

However, although the price range does capture some of the uncertainty involved with model choice, it does not necessarily capture the uncertainty in a suitable way. First, the price range measure may fail to detect model uncertainty, for example when prices are equal across pricing measures, but the hedging strategies differ. In other words, $\mu(X) = 0$ does not imply the absence of model uncertainty. To illustrate this, consider for example a Black-Scholes type model with time-dependent volatility, that is, the asset price dynamics are given by

$$\mathrm{d}S_t = S_t \left(r\,\mathrm{d}t + \sigma(t)\,\mathrm{d}W_t \right), \tag{26}$$

with r the risk-free interest rate and $W = (W_t)_{t\geq 0}$ a standard Brownian motion. The calibration condition to match a traded call option with maturity T is

$$\frac{1}{T}\int_0^T \sigma(t)^2 \mathrm{d}t = \Sigma^2, \tag{27}$$

where Σ is the implied Black-Scholes volatility of the benchmark call option. Now let $\sigma(t) = at+b$ with $t \in [0,T]$ be a line with $a < 0$ to reproduce commonly observed volatility surfaces with high short-term volatility. The calibration condition then yields $b = 1/6(\sqrt{36\Sigma^2 - 3a^2T^2} - 3aT)$, and $b \in \mathbb{R}$ if $a \geq -(2\sqrt{3}\Sigma)/T$. Now suppose that the model set $\mathcal{Q}$ is represented by different choices of b in some range $[b_{\min}, b_{\max}]$. For any European payoff X with maturity T we have $\mu(X) = 0$, if the price of X depends only on the distribution of the terminal value S_T. At the same time, the prices and hedge ratios at $T/2$ will differ across the models because the remaining volatilities differ.

Second, the price range measure is incompatible with (regulatory) capital charges for other risk types. For example, market risk is typically quantified in terms of value-at-risk, which is the quantile of the loss distribution at an e.g. 95% or 99% level. If model uncertainty is measured in terms of the price range, then the risks may possibly be traded off against each other, as they do not measure the same loss quantities.

Finally, in Section 4.1 it is shown that price differences in two models are expected P&L from hedging a claim in one model relative to another model. The actual loss from model uncertainty can be much larger than the expected loss. A value-at-risk or expected-shortfall based approach at a 95%- or 99%-confidence level is a more conservative approach for creating a reserve buffer.

B Construction of extended probability space

We recall the definition of a regular conditional probability:

Definition 18 (Shiryaev (1996), Section II.7)**.** A function $\overline{\mathbb{P}}(\omega; B)$, defined for all $\omega \in \Omega$ and $B \in \mathcal{F}$, is a regular conditional probability with respect to $\mathcal{G}$ if

(a) $\overline{\mathbb{P}}(\omega; \cdot)$ is a probability measure on $\mathcal{F}$ for every $\omega \in \Omega$;
(b) For each $B \in \mathcal{F}$ the function $\overline{\mathbb{P}}(\omega; B)$, as a function of ω, is a variant of the conditional probability $\overline{\mathbb{P}}(B|\mathcal{G})(\omega)$, i.e. $\overline{\mathbb{P}}(\omega; B) = \overline{\mathbb{P}}(B|\mathcal{G})(\omega)$ (a.s.).

For the construction of an extended probability space (see Section 4.3) we start with a set of probability measures $\mathcal{P}$ on $(\Omega, \mathcal{F})$. Let $(\tilde{\Omega}, \tilde{\mathcal{F}})$, $(\Theta, \boldsymbol{\Theta})$ be measurable spaces and let $\theta : \tilde{\Omega} \to \Theta$ be a measurable mapping such that $\theta \mapsto \overline{\mathbb{P}}_\theta$ is a measurable mapping from $(\Theta, \boldsymbol{\Theta})$ to $(\mathcal{P}, \sigma(\mathcal{P}))$. Without loss of generality we may assume that $\Theta \subseteq \mathbb{R}$. We wish to construct a probability measure $\overline{\mathbb{P}}$ on $(\Omega \times \tilde{\Omega}, \mathcal{F} \times \tilde{\mathcal{F}})$ so that $\overline{\mathbb{P}}_\theta = \overline{\mathbb{P}}(\cdot|\theta)$ holds $\overline{\mathbb{P}}$–a.s., that is, where the partial averaging property that defines conditional expectation is fulfilled:

$$\int_{\theta \in A} \overline{\mathbb{P}}_\theta(B) \, \mathrm{d}\overline{\mathbb{P}} = \int_{\theta \in A} \mathbf{1}_B \, \mathrm{d}\overline{\mathbb{P}}, \quad A \in \mathcal{B}(\Theta), \quad B \in \mathcal{F}. \tag{28}$$

Rewrite the left-hand side as

$$\int_{\theta \in A} \overline{\mathbb{P}}_\theta(B) \, \mathrm{d}\overline{\mathbb{P}} = \int \overline{\mathbb{P}}_{\theta(\tilde{\omega})}(B) \, \mathbf{1}_{\{\theta(\tilde{\omega}) \in A\}} \, \overline{\mathbb{P}}(\mathrm{d}\tilde{\omega}) = \int \overline{\mathbb{P}}_a(B) \, \mathbf{1}_{\{a \in A\}} \, \mu(\mathrm{d}a), \tag{29}$$

where μ is the distribution of θ on $(\tilde{\Omega}, \tilde{\mathcal{F}})$. Similarly, we obtain for the right-hand side,

$$\int_{\theta \in A} \mathbf{1}_B \, \mathrm{d}\overline{\mathbb{P}} = \int \mathbf{1}_{\{B, \theta(\tilde{\omega}) \in A\}} \, \overline{\mathbb{P}}(\mathrm{d}\tilde{\omega}) = \overline{\mathbb{P}}(B, \theta \in A). \tag{30}$$

If we *define* $\overline{\mathbb{P}}$ in this way, that is, as

$$\int \overline{\mathbb{P}}_a(B) \, \mathbf{1}_{\{a \in A\}} \, \mu(\mathrm{d}a) =: \overline{\mathbb{P}}(B, \theta \in A), \tag{31}$$

then it remains to show that $\overline{\mathbb{P}}$ is indeed a probability measure. But this is easily established since for all $a \in \Theta$, $\overline{\mathbb{P}}_a$ are probability measures. By construction $\mathcal{P}$ is a regular conditional probability on $(\Omega \times \tilde{\Omega}, \mathcal{F} \times \tilde{\mathcal{F}})$.

C Proof of Lemma 15

Proof. Axiom 1: The ($\mathcal{F}_T$-measurable) benchmark instrument H_i can be statically hedged, that is, $u_i = 1$, $u_j = 0$, $j = 1, \ldots, I$ and $j \neq i$, and $\phi = (0)_{0 \le t \le T}$. The associated loss is $L_t(H_i, \Phi) = 0$ $\overline{\mathbb{P}}$–a.s., so that $\mu_t^{\mathbb{Q}}(H_i) = 0$.

Axiom 2: Let $X \in \mathcal{C}$ a claim such that

$$\hat{\mathbb{Q}}\left(\mathbb{E}[X|\mathcal{F}_t] = \mathbb{E}[X|\mathcal{F}_0] + \int_0^t \phi_t(X) \, \mathrm{d}S_t\right) = 1, \forall \hat{\mathbb{Q}} \in \mathcal{Q}, \forall t. \tag{32}$$

Choose $\Phi = (\phi(X), 0, \ldots, 0)$. Then clearly $L_t(X, \Phi) = 0, \overline{\mathbb{P}}$–a.s. and thus $f(L_t(X, \Phi)) = 0$ implying $\mu_t^{\mathbb{Q}}(X) = 0$ since $0 \le \mu_t^{\mathbb{Q}}(X) \le f(L_t(X, \Phi))$.

Let now $\tilde{X} \in \mathcal{C}$. The loss is a linear function of the hedging strategy by Assumption 3, so that $\overline{\mathbb{Q}}_a(L_t(\tilde{X} + X, \tilde{\Phi} + \Phi) = L_t(\tilde{X}, \tilde{\Phi})) = 1$, for all $a \in \Theta$ and for all $\tilde{\Phi} \in \Pi(\tilde{X})$. Therefore, $f(L_t(\tilde{X} + X, \tilde{\Phi} + \Phi)) = f(L_t(\tilde{X}, \tilde{\Phi})$ by assumption on f. Since $\{\tilde{\Phi} + \Phi | \tilde{\Phi} \in \Pi(\tilde{X})\} \subseteq \Pi(\tilde{X} + X)$ we obtain $\mu_t^{\mathbb{Q}}(\tilde{X} + X) \le \mu_t^{\mathbb{Q}}(\tilde{X})$.

If X is such that (32) holds, then (32) also holds with X replaced by $-X$ due to the linearity property of Assumption 3. Then the same argument applied to $((\tilde{X} + X) - X)$ proves that $\mu_t^{\mathbb{Q}}(X + \tilde{X}) = \mu_t^{\mathbb{Q}}(\tilde{X})$.

Axiom 3: Assume that $\mu_t^{\mathbb{Q}}$ is not convex, so there exist $\lambda \in [0, 1]$ and claims $X, \tilde{X} \in \mathcal{C}$ such that

$$\mu_t^{\mathbb{Q}}(\lambda X + (1 - \lambda)\tilde{X}) > \lambda \mu_t^{\mathbb{Q}}(X) + (1 - \lambda)\mu_t^{\mathbb{Q}}(\tilde{X}). \tag{33}$$

Then, by definition of $\mu_t^{\mathbb{Q}}$, there exist strategies $\Phi \in \Pi(X)$ and $\tilde{\Phi} \in \Pi(\tilde{X})$ such that

$$\inf_{\hat{\Phi} \in \Pi(\lambda X + (1-\lambda)\tilde{X})} L_t(\lambda X + (1 - \lambda)\tilde{X}, \hat{\Phi}) > \lambda L_t(X, \Phi) + (1 - \lambda) L_t(\tilde{X}, \tilde{\Phi}). \tag{34}$$

However, again by Assumption 3, $\lambda\Phi + (1 - \lambda)\tilde{\Phi} \in \Pi(\lambda X + (1 - \lambda)\tilde{X})$ and

$$L_t(\lambda X + (1 - \lambda)\tilde{X}, \lambda\Phi + (1 - \lambda)\tilde{\Phi}) = \lambda L_t(X, \Phi) + (1 - \lambda) L_t(\tilde{X}, \tilde{\Phi}), \tag{35}$$

and since f is convex

$$
\begin{aligned}
f(L_t(\lambda X + (1-\lambda)\tilde{X}, \lambda\Phi + (1-\lambda)\tilde{\Phi})) &= f(\lambda L_t(X,\Phi) + (1-\lambda)L_t(\tilde{X},\tilde{\Phi})) \\
&\leq \lambda f(L_t(X,\Phi)) + (1-\lambda) f(L_t(\tilde{X},\tilde{\Phi})). \qquad (36)
\end{aligned}
$$

This in turn implies that

$$
\inf_{\hat{\Phi}\in\Pi(\lambda X+(1-\lambda)\tilde{X})} f(L_t(\lambda X + (1-\lambda)\tilde{X}, \hat{\Phi})) \leq \lambda f(L_t(X,\Phi)) + (1-\lambda) f(L_t(\tilde{X},\tilde{\Phi})), \qquad (37)
$$

which contradicts Equation (34).

Axiom 4: Let $X \in \mathcal{C}$ and $(u_1, \ldots, u_I)$ be given. For any $\Phi \in \Pi(X)$, we have $\Phi + \tilde{\Phi} \in \Pi(X + \sum_I u_i H_i)$, where $\tilde{\Phi} = ((0)_{0\leq s\leq T}, u_1, \ldots, u_I)$, and furthermore $L_t(X + \sum_I u_i H_i, \Phi + \tilde{\Phi}) = L_t(X, \Phi)$, so that

$$
f(L_t(X + \textstyle\sum_{i=1}^{I} u_i H_i, \Phi + \tilde{\Phi})) = f(L_t(X,\Phi)). \qquad (38)
$$

From $\{\tilde{\Phi} + \Phi | \Phi \in \Pi(X)\} \subseteq \Pi(X + \sum_{i=1}^{I} u_i H_i)$ we obtain $\mu_t^{\mathbb{Q}}(X + \sum_{i=1}^{I} u_i H_i) \leq \mu_t^{\mathbb{Q}}(X)$. □

References

H. Ahn and P. Wilmott. Dynamic hedging is dead! Long live static hedging! *Wilmott magazine*, pages 80–87, January 2008.

H. Akaike. Information theory and an extension of the maximum likelihood principle. In B. N. Petrov and F. Csaki, editors, *Second International Symposium on Information Theory*, pages 267–281, Budapest, 1973. Akadémiai Kiado.

H. Albrecher and P. Mayer. Semi-static hedging strategies for exotic options. In R. Kiesel, M. Scherer, and R. Zagst, editors, *Alternative Investments and Strategies*, pages 345–373. World Scientific, Singapore, 2010.

P. Artzner, F. Delbaen, J. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999.

M. Ascheberg, B. Bick, and H. Kraft. Hedging structured credit products during the credit crisis: A horse race of 10 models. *Journal of Banking & Finance*, 37(5):1687–1705, 2013.

K. Bannör and M. Scherer. Capturing parameter uncertainty with convex risk measures. *European Actuarial Journal*, 3:97–132, 2013.

BIS. Revisions to the Basel II market risk framework. Basel Committee on Banking Supervision, Bank for International Settlements, February 2011.

A. N. Borodin and P. Salminen. *Handbook of Brownian motion: facts and formulae.* Springer, 2002.

C. M. Boucher, J. Danielsson, P. S. Kouontchou, and B. B. Maillet. Risk Model-at-Risk. *Journal of Banking & Finance*, 44:72–92, 2014. ISSN 0378-4266.

N. Boyarchenko, M. Cerrato, J. Crosby, and S. Hodges. No good deals – no bad models. Working paper, January 2014.

N. Branger, E. Krautheim, C. Schlag, and N. Seeger. Hedging under model misspecification: All risk factors are equal, but some are more equal than others... *Journal of Futures Markets*, 32(5):397–430, 2012.

D. T. Breeden and R. H. Litzenberger. Prices of state-contingent claims implicit in option prices. *The Journal of Business*, 51(4):621–651, 1978.

T. Breuer and I. Csiszar. Measuring distribution model risk. *Mathematical Finance*, (forthcoming), 2014.

D. Brigo and F. Mercurio. Lognormal-mixture dynamics and calibration to market volatility smiles. *International Journal of Theoretical and Applied Finance*, 5(04):427–446, 2002.

K. Burnham and D. Anderson. *Model selection and multimodel inference: A practical information-theoretic approach*. Springer, New York, 2nd edition, July 2002. ISBN 0387953647. URL http://www.worldcat.org/isbn/0387953647.

K. Burnham and D. Anderson. Multimodel inference – Understanding AIC and BIC in model selection. *Sociological Methods & Research*, 33(2):261–304, 2004.

P. Carr and D. Madan. Towards a theory of volatility trading. In R. Jarrow, editor, *Risk Book on Volatility*. Risk, New York, 1998. pp. 417–427.

P. Carr and L. Wu. Variance risk premiums. *Review of Financial Studies*, 22(3):1311–1341, 2009.

P. Carr, K. Ellis, and V. Gupta. Static hedging of exotic options. *The Journal of Finance*, 53 (3):1165–1190, 1998.

J. E. Cavanaugh. Unifying the derivations for the Akaike and corrected Akaike information criteria. *Statistics & Probability Letters*, 33(2):201–208, 1997.

P. Cohort, P.-E. L. D. Vehel, and F. Patras. Toward model value-at-risk: bespoke CDO tranches, a case study. *Journal of Risk Model Validation*, 7(3):21–34, 2013.

R. Cont. Model uncertainty and its impact on the pricing of derivative instruments. *Mathematical Finance*, 16(3):519–547, 2006.

R. Cont and C. Jessen. Constant proportion debt obligations (CPDOs): modeling and risk analysis. *Quantitative Finance*, 12(8):1199–1218, 2012.

R. Cont and P. Tankov. *Financial Modelling with Jump Processes*. Chapman & Hall/CRC, 2004.

M. Crouhy, D. Galai, and R. Mark. Model risk. *Journal of Financial Engineering*, 7:267–288, 1998.

M. Denault. Coherent allocation of risk capital. *Journal of Risk*, 4:1–34, 2001.

L. Denis and C. Martini. A theoretical framework for the pricing of contingent claims in the presence of model uncertainty. *The Annals of Applied Probability*, pages 827–852, 2006.

E. Derman. Model risk. *Risk Magazine*, 9(5), May 1996.

N. Detering and N. Packham. Model risk in incomplete markets with jumps. In K. Glau, M. Scherer, and R. Zagst, editors, *Innovations in Risk Management*, volume forthcoming of *Springer Proceedings in Mathematics & Statistics*, 2014.

B. Dupire. Pricing with a smile. *Risk*, 7(1):18–20, 1994.

EBA. Discussion paper on draft regulatory technical standards on prudent valuation, under Article 100 of the draft Capital Requirements Regulation (CRR). Discussion Paper, European Banking Authority, November 2012.

N. El Karoui, M. Jeanblanc-Picqué, and S. Shreve. Robustness of the Black and Scholes formula. *Mathematical Finance*, 8(2):93–126, 1998.

A. Elices and E. Giménez. Applying hedging strategies to estimate model risk and provision calculation. *Quantitative Finance*, 13(7):1015–1028, 2013.

Federal Reserve. Supervisory guidance on model risk management. Board of Governors of the Federal Reserve System, Office of the Comptroller of the Currency, SR Letter 11-7 Attachment, April 2011.

S. Figlewski. Derivatives risks, old and new. *Brookings-Wharton Papers on Financial Services*, 1:159–217, 1998.

H. Föllmer and A. Schied. Convex measures of risk and trading constraints. *Finance and Stochastics*, 6(4):429–447, 2002.

M. Frittelli and E. Rosazza Gianin. Putting order in risk measures. *Journal of Banking & Finance*, 26(7):1473–1486, 2002.

D. Galai. Tests of market efficiency of the Chicago Board Options Exchange. *The Journal of Business*, 50(2):167–197, 1977.

J. Gatheral. *The Volatility Surface*. Wiley, Hoboken, NJ, 2006.

P. Glasserman. *Monte Carlo Methods in Financial Engineering*. Springer, 2004.

P. Glasserman and X. Xu. Robust risk measurement and model risk. *Quantitative Finance*, 14 (1):29–58, 2014.

M. Gordy and S. Willemann. Constant proportion debt obligations: A postmortem analysis of rating models. *Management Science*, 58(3):476–492, 2012.

C. Gourieroux and A. Monfort. *Statistics and Econometric Models*, volume 2. Cambridge University Press, 1995.

E. Heitfield. Parameter uncertainty and the credit risk of collateralized debt obligations. Working paper, Federal Reserve Board, 2008.

P. Hénaff and C. Martini. Model validation: theory, practice and perspectives. *Journal of Model Risk Validation*, 5(4):3–15, 2011.

S. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2):327–343, 1993. ISSN 0893-9454.

J. Hoeting, D. Madigan, A. Raftery, and C. Volinsky. Bayesian model averaging: A tutorial. *Stat. Sci.*, 14(4):382–417, 1999.

J. Hull and W. Suo. A methodology for assessing model risk and its application to the implied volatility function model. *Journal of Financial and Quantitative Analysis*, 37(2):297–318, 2002.

C. M. Hurvich and C.-L. Tsai. Regression and time series model selection in small samples. *Biometrika*, 76(2):297–307, 1989.

R. Jarrow. Hedging derivatives with model error. *Quantitative Finance*, 12(6):855–863, 2012.

P. Jorion. Risk management lessons from the credit crisis. *European Financial Management*, 15(5):923–933, 2009.

M. Kalkbrener. An axiomatic approach to capital allocation. *Mathematical Finance*, 15(3): 425–437, 2005.

J. Kerkhof, B. Melenberg, and H. Schumacher. Model risk and capital reserves. *Journal of Banking & Finance*, 34(1):267–279, 2010.

F. H. Knight. *Risk, uncertainty and profit.* Houghton Mifflin, Boston, 1921.

A. McNeil, R. Frey, and P. Embrechts. *Quantitative Risk Management.* Princeton University Press, Princeton, NJ, 2005.

R. C. Merton. Option pricing when the underlying stock returns are discontinuous. *Journal of Financial Economics*, 3:125–144, 1976.

R. Merton, M. Scholes, and M. Gladstein. The returns and risk of alternative call option portfolio investment strategies. *Journal of Business*, pages 183–242, 1978.

R. Merton, M. Scholes, and M. Gladstein. The returns and risks of alternative put-option portfolio investment strategies. *Journal of Business*, pages 1–55, 1982.

M. Morini. *Understanding and Managing Model Risk: A Practical Guide for Quants, Traders and Validators.* Wiley, Hoboken, NJ, 2011.

M. Nalholm and R. Poulsen. Static hedging and model risk for barrier options. *Journal of Futures Markets*, 26(5):449–463, 2006.

N. Packham, L. Schloegl, and W. Schmidt. Credit gap risk in a first passage time model with jumps. *Quantitative Finance*, 13(12):1871–1889, 2013.

K. Sato. *Lévy Processes and Infinitely Divisible Distributions.* Cambridge University Press, 1999.

W. Schoutens. *Lévy Processes in Finance.* John Wiley & Sons, 2003.

W. Schoutens, E. Simons, and J. Tistaert. A perfect calibration! now what? *Wilmott magazine*, 2:66–78, 2004.

G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464, 1978.

A. N. Shiryaev. *Probability.* Springer, Berlin, 2nd edition, 1996.

M. Soner, N. Touzi, and J. Zhang. Quasi-sure stochastic analysis through aggregation. *Electronic Journal of Probability*, 16:1844–1879, 2011. Article number 67.

K. Takeuchi. Distribution of informational statistics and a criterion of model fitting (in Japanese). *Suri-Kagaku (Mathematical Sciences)*, 153(1):12–18, 1976.

D. van Deventer. CDOs and the credit crisis: Complexity and model risk in the collateralized debt obligation market are severe. Working Paper, Kamakura Corporation, April 2008.